



Contribution ID: 22

Type: **not specified**

Utilizing tools made for "Big Data" to analyse Ftrace data - making it fast and easy

Tuesday, 10 September 2019 15:45 (45 minutes)

Tools based on low level tracing tend to generate large amounts of data, typically outputted in some kind of text or binary format. On the other hand the predefined data analysis features of those tools are often useless when it comes to solving a nontrivial or very user-specific problem. This is when the possibility to make sophisticated analysis via scripting can be extremely useful.

Fast and easy scripting inside the tracing data is possible if we take advantage of the already existing infrastructure, originally developed for the purposes of the "Big Data" and ML industries. A PoC interface for accessing Ftrace data in Python (via NumPy arrays) will be demonstrated, together with few examples of analysis scripts. Currently the prototype of the interface is implemented as an extension of KernelShark. This is a work in progress, and we hope to receive advice from experts in the field to make sure the end result works seamlessly for them.

I agree to abide by the anti-harassment policy

I confirm that I am already registered for LPC 2019

Primary author: KARADZHOV, Yordan (VMware)

Presenter: KARADZHOV, Yordan (VMware)

Session Classification: LPC Refereed Track