



Contribution ID: 237

Type: **not specified**

Developing eBPF profiler for polyglot cloud-native applications

Monday, 12 September 2022 17:00 (30 minutes)

One of the important jobs of system-wide profilers is to capture stack traces without requiring recompilation or redeployment of profiled applications. This becomes difficult when the profiler has to deal with the binaries compiled in different languages. Heavy lifting for the stack unwinding is done by the kernel if frame pointers are present or if the binary has ORC - in kernel debug information format information available. Although most modern compilers have an option to scrap the frame pointers for performance gain.

In this talk we will talk about how we are experimenting with using eBPF to extend the existing stack unwinding facility in the Linux kernel. We will discuss how we are walking the stacks of interpreted languages, such as Ruby, as well as runtimes with JITs, like the JVM. And how extending the current stack unwinding facility can be useful for such cases.

I agree to abide by the anti-harassment policy

Yes

Primary authors: THAKKAR, Vaishali; HONDUVILLA COTO, Javier

Presenters: THAKKAR, Vaishali; HONDUVILLA COTO, Javier

Session Classification: eBPF & Networking

Track Classification: eBPF & Networking Track