



Contribution ID: 326

Type: **not specified**

bpftime: Fast uprobes with user space BPF runtime

Tuesday, 14 November 2023 09:00 (30 minutes)

In kernel operations, the uprobe component of eBPF often faces performance inefficiencies, primarily due to the overheads introduced by context switches. Transitioning to userspace, eBPF can bypass these context switch-induced delays, leading to optimized performance. Moreover, this transition facilitates greater configurability without requiring root access or privileges, thus reducing the kernel attack surface.

In this talk, we will introduce bpftime, a prototype userspace eBPF runtime. It offers rapid uprobe and syscall hook capabilities: userspace uprobes can be 10x faster than kernel uprobes, without the necessity of two context switches. It can also programmatically hook syscalls of a process safely and efficiently.

Utilizing binary rewriting, we enable uprobe and syscall hooks: These can trace or patch the execution of a function, and hook, filter, or redirect all syscalls of a process with an eBPF userspace runtime. This runtime can be injected into any running process without the need for a restart or manual recompilation.

Additionally, we have implemented interprocess eBPF Maps in shared userspace memory for summary aggregation or control plane communication. Compatibility is ensured with existing eBPF toolchains like clang and libbpf for developing userspace eBPF without any modifications. We support CO-RE via BTF, and offer userspace host function access, broadening the utility and ease of use of the bpftime runtime.

Primary author: ZHENG, Yusheng (PLCT Lab)

Co-authors: YU, Tong (PLCT Lab); YANG, Yiwei (UCSC)

Presenter: ZHENG, Yusheng (PLCT Lab)

Session Classification: eBPF & Networking

Track Classification: eBPF & Networking Track