

BPF user experience rough edges

Thursday, September 23, 2021 7:50 AM (40 minutes)

This talk highlights a few rough edges in the overall BPF user experience that we have observed while building services with BPF at Cloudflare. We will showcase a set of problems, analyze their cause, and present possible workarounds. The goal of the talk is to share collected know-how with other users, and trigger discussions on potential improvements.

Collected cases fall into two distinct categories:

1. issues when running BPF with as few capabilities as possible,
2. issues when loading generated BPF programs.

Within the first group we are going to cover such topics as:

- locked memory limit (still relevant because present in LTS kernels),
- credentials control on BPF links,
- access control on BPF maps,
- accessing pinned objects under `/sys/fs/bpf`,
- incompatibility between existing socket maps.

In the second category, we'll cover various clang / LLVM optimizations that cause generated C to fail with only small input changes:

- optimized out packet bounds checks,
- stack spilling,
- register "mirroring", where clang thinks they have the same value but not the verifier,
- inter generated code optimizations.

We'll also discuss how we're switching to a hybrid static C & generated eBPF model, and fuzzing the eBPF generator.

I agree to abide by the anti-harassment policy

I agree

Primary authors: SITNICKI, Jakub (Cloudflare); FABRE, Arthur (Cloudflare)

Presenters: SITNICKI, Jakub (Cloudflare); FABRE, Arthur (Cloudflare)

Session Classification: BPF & Networking Summit

Track Classification: Networking & BPF Summit (Closed)