

Watching the super powers

Tuesday, 21 September 2021 07:50 (40 minutes)

BPF programs are critical system components performing core networking functionality, system audit logs, tracing, and runtime security enforcement to list a few. Charged with such crucial tasks, how do we audit the BPF subsystem itself to ensure system bugs are noticed and malicious attackers can not subtly manipulate these components, inject new programs, or quietly run their own BPF programs?

One proposal is to sign BPF programs following the model used for years to sign kernel modules. In this model the BPF programs loaded into the kernel are signed and verified to ensure only authorized programs can be loaded. Although we support such efforts we believe they are insufficient to actually provide meaningful security guarantees. Unlike kernel modules of old BPF programs tend to have a control plane that are tightly coupled with the application

where signing the BPF program only covers the most obvious attacks.

In this talk, using real-world examples, we show signing BPF programs provides only minimal improvement over the current model. Instead we propose a robust audit system and enforcement of BPF system calls to ensure access to critical control paths and enforcing loaders of said programs are known. Further, by scanning programs at load time we can do fine grained permissions, e.g. only allowing specific BPF helpers and maps to be exposed in the targeted file systems. Finally, by doing runtime auditing and enforcement we can provide fine grained per user policies based on the trust worthiness of the user. How do we propose to build such a platform? Using BPF of course! By loading a core set of BPF programs in early boot we will show how to implement the proposed model.

I agree to abide by the anti-harassment policy

I agree

Primary author: FASTABEND, John (Isovalent)

Presenter: FASTABEND, John (Isovalent)

Session Classification: BPF & Networking Summit

Track Classification: Networking & BPF Summit (Closed)