Contribution ID: **130** Type: **not specified**

bpfilter - BPF based firewall

Wednesday, September 22, 2021 10:20 AM (40 minutes)

Motivation

Iptables has become a synonym of a firewall in Linux world. Although there is a nftables which is supposed to replace iptables, iptables will exist for decades more because of its popularity and ubiquity.

With the growing widespread use of BPF technology and its benefits there is a temptation to apply the technology for the firewalling purposes.

Problem Statement

Despite its advantages iptables is also known for its dark side - performance and security related issues. What if it's possible to keep the iptables' ABI and replace its implementation with something more performant and secure by nature?

Such an approach will keep the existing solutions to work and remove an overhead of switching to a new technology.

Approach

There was a RFC patchset back in 2018 which proposed a BPF based firewall - bpfilter. From a bird's eye view bpfilter is a compiler implemented as a user mode helper kernel module. bpfilter analyses an iptables' ruleset and synthesizes an equivalent BPF program. When bpfilter kernel module is loaded it starts a userspace process that has an IPC with its kernelspace part. Most of the bpfilter functionaility is implemented in the userspace process what significantly simplifies its development and improves security. The kernel part hooks into the kernel iptables ABI and transparently for the userspace consumer passes control to the userspace process. bpfilter userspace process "compiles" iptables' ruleset into a BPF program and passes control back to the kernel. This approach allows to transparently replace iptables' implementation without breaking its consumers and gain all the benefits of BPF ecosystem.

Results

While the initial patchset was abandoned in 2021 there was an attempt to ressurect the patchset. Two versions of the updated patchset were submitted to the bpf@ mailing list and the third iteration is in the process of preparation. Currently bpfilter is able to process basic rules in INPUT and OUTPUT chains and translate them into equivalent XDP and TC programs. bpfilter has an easy way to add new matches and extensions in terms of iptables.

Conclusions

The idea to treat a firewall as a compiler is seductive - as such an approach provides more opportunities for performance optimisations due to a more precise context. Combining it with the existing BPF performance and security features and putting on top of it its userspace nature - this might sound as the next firewall for Linux.

I agree to abide by the anti-harassment policy

I agree

Primary author: BANSHCHIKOV, Dmitrii (Facebook)

Presenter: BANSHCHIKOV, Dmitrii (Facebook)

Session Classification: BPF & Networking Summit

Track Classification: Networking & BPF Summit (Closed)