


Agilio® eBPF for Fast, Versatile Security

HIGH PERFORMANCE, FLEXIBLE PROTECTION FOR CLOUD DATA CENTER AND EDGE COMPUTING APPLICATIONS

A SHIFT IN THE LINUX COMMUNITY

The Linux community is adopting a new way to protect network traffic. Rather than renovating iptables or replacing iptables with nftables, moving to bpf brings the promise of stronger, faster and more flexible network security. Although Berkeley Packet Filter (BPF) is not new, its incorporation into the Linux kernel and its subsequent extension (eBPF) promote not just programmable network filtering, but new and secure flexibility to many parts of the Linux kernel. This is what makes it such an important shift.



NETWORK AGILITY

eBPF uses a byte code language that is just-in-time (JIT) compiled to execute programs efficiently in the kernel. Every program written in user space is first verified for safe execution before it is allowed to execute on the kernel VM. Because programs can be associated with a number of kernel attach points, system designers have tremendous flexibility in adding functionality without introducing new security concerns or sacrificing performance. For network filtering, bpf is a faster and more secure path to firewalling in the kernel; faster than iptables or nftables. But, if adopting this new path is not attractive, the eXpress Data Path (XDP) or Traffic Control (TC) are good options. eBPF is a clean, safe and efficient architecture for extending kernel control.



EBPF PROPERTIES

- Protected kernel execution
- Static program verification
- Programs guaranteed to terminate

EBPF FUNCTIONS

- XDP offload
- TC offload
- Match/action
- Filtering
- Load balancing
- DDoS mitigation
- Maps for stateful programs
- Chained filter functions

BENEFITS

The value of having eBPF in the Linux kernel should not be underestimated. It permits service designers and appliance builders the ability to write programs for differentiated services and inject them into the Linux kernel without compromising kernel integrity or performance.

These programs can remain proprietary because they do not violate kernel licensing terms while simultaneously extending kernel functionality. If you are using a supported kernel, they would not violate support contract terms either. In addition, with network hardware acceleration some of these services can be implemented at wire speed without impact to application performance.


In clouds and data centers, this can translate into low cost, high speed, customized and agile firewall policies at the edge, core or end servers.

For appliance builders, this enables exclusive extensions to the Linux kernel for rapid innovation and with transparent hardware offload, the same code used in low-end appliance models can instantly apply to high-end models.

SIMPLE COMPARISON

To illustrate the difference between filtering technologies, a test that highlights the advantages of bpfILTER with and without hardware acceleration was performed. The test was a simple packet drop with packets arriving at line rate on the following setup:

- Intel® Xeon® CPU E5-2630 v3 @ 2.40 GHz; single CPU, 8 cores 16 threads
- Netronome Agilio CX, 1x40GbE SmartNIC
- Linux kernel v4.16
- One single rule for iptables and nftables
- 64 byte data packets



What becomes obvious is the remarkable performance improvement between bpfILTER, iptables and nftables; a 3.5X improvement over iptables and a 3.1X improvement over nftables (~6X and ~5X respectively with hardware offload). More importantly, the test with transparent hardware offload provided line rate performance with zero impact to the host CPU, as the SmartNIC performed the processing needed to drop packets.

THE NETRONOME ADVANTAGE

Because of the eBPF design, Netronome's processors are uniquely positioned to execute eBPF programs in hardware through the use of our own JIT compiler that transparently translates eBPF code into NFP machine code, greatly accelerating the rules written for the kernel by running them directly on the SmartNIC. Netronome's performance enhancements are exceptional and could not be easier to use because they are included in the Linux kernel.

NETRONOME

Netronome Systems, Inc.
 2903 Bunker Hill Lane, Suite 150
 Santa Clara, CA 95054
 Tel: 408.496.0022 | Fax: 408.586.0002
www.netronome.com

©2018 Netronome. All rights reserved.

Netronome, the Netronome logo, and Agilio are trademarks or registered trademarks of Netronome Systems, Inc. All other trademarks mentioned are registered trademarks or trademarks of their respective owners in the United States and other countries.