



Contribution ID: 294

Type: **not specified**

A BPF map for online packet classification

Tuesday, 13 September 2022 11:00 (30 minutes)

There is a growing need in online packet classification for BPF-based networking solutions. In particular, in cilium we have two use cases: the PCAP recorder for the standalone XDP load balancer [1] and the k8s network policies. The PCAP recorder implementation suffers from slow and dangerous updates due to runtime recompilation, and both use cases require specifying port ranges in rules, which is currently not supported.

At the moment there are two competing algorithms for online packet classification: Tuple Merge [2] and Partition Sort [3]. The Tuple Merge algorithm is using hash tables to store rules, and the Partition Sort is using multi-dimensional Interval trees. Thus, both of algorithms are [nearly?] impossible to implement in “pure” BPF due to lack of functionality and also due to verifier complexity limits.

We propose a new BPF map for packet classification and an API which can be used to adapt this map to different practical use cases. The map is not tied to the use of a specific algorithm, so any of brute force, tuple merge, partition sort or a future state-of-art algorithm can be used.

[1] <https://cilium.io/blog/2021/05/20/cilium-110/#pcap>

[2] <https://nonsns.github.io/paper/rossi19ton.pdf>

[3] https://www.cse.msu.edu/~yingchar/PDF/sorted_partitioning_icnp2016.pdf

I agree to abide by the anti-harassment policy

Yes

Primary author: PROTOPOPOV, Anton (Isovalent)

Presenter: PROTOPOPOV, Anton (Isovalent)

Session Classification: eBPF & Networking

Track Classification: eBPF & Networking Track