

Packet mark in the Cloud Native world

Monday, August 24, 2020 7:45 AM (45 minutes)

The 32-bit “mark” associated with the skb has served as a metadata exchange format for Linux networking subsystems since the beginning of the century. Over that time, the interpretation and reuse of the field has grown to encapsulate a wide range of networking use cases, expanding to touch everything from iptables, tc, xfrm, openvswitch, sockets, routing, to eBPF. In recent years, more than a dozen network control applications have been written in the Cloud Native space alone, many of which are using the packet mark in different ways to solve networking problems. The kernel facilities define no specific semantics to these bits, which leaves it up to these applications to co-ordinate to avoid incompatible mark usage.

This talk will explore use cases for sharing metadata between Linux subsystems in light of recent containerization trends, including but not limited to: application identity, firewalling, ip masquerade, network isolation, service proxying and transparent encryption. Beyond that, Cilium’s particular usage will be discussed with approaches used to mitigate conflicts due to the inevitable overload of the mark.

I agree to abide by the anti-harassment policy

I agree

Primary author: STRINGER, Joe (Cilium.io)

Presenter: STRINGER, Joe (Cilium.io)

Session Classification: Networking and BPF Summit

Track Classification: Networking & BPF Summit