

Linux Plumbers Conference 2024



Contribution ID: 391

Type: **not specified**

What makes the panda sad in the Linux network stack today?

Wednesday, 18 September 2024 12:30 (30 minutes)

Three years ago, we had the pleasure of giving a talk at Linux Plumbers about rough edges in BPF user experience. Attendees might recall that we found quite a few reasons for the panda to be sad about BPF UX back then.

This time, we would like to come back and present an assortment of snags we have encountered in the Linux network stack itself:

- Is it possible to have a proper loopback subnet like 127/8 for IPv6?
- How TS.Recent TCP timestamp resolution can lead to port exhaustion?
- UDP segmentation offload does wonders for throughput, but can you always use it?
- Why sourcing return traffic when using BPF socket lookup is tricky for UDP?
- How early demux can get in the way of forwarding traffic?
- Why IP_BIND_ADDRESS_NO_PORT does not work the way you think for UDP?
- Do listen() semantics for UDP make sense when we consider QUIC?

Some of these problems we have worked on and managed to solve in collaboration with the Linux upstream community. Some we have only discussed publicly but intend to propose and contribute a solution for, while others we just acknowledge and present a workaround recipe, if one exists.

We hope to engage the audience to learn who else has run into any of the presented obstacles, are there any alternative approaches that we have not considered, and to collect input on how to best solve them.

Primary author: SITNICKI, Jakub (Cloudflare)

Presenter: SITNICKI, Jakub (Cloudflare)

Session Classification: Networking Track

Track Classification: Networking Track