Rich Packet Metadata: The Saga Continues

Jakub Sitnicki
Cloudflare

FOSDEM 2026
Brussels, Belgium

Agenda

0 Motivation & use cases

1 Part One: SKB Traits

2 Part Two: XDP Metadata

3 Part Three: SKB Extension

Cloudflare’s goal to hire 1,111 interns in 2026

https://blog.cloudflare.com/cloudflare-1111-intern-program/

Motivation & use cases for packet metadata

Motivation – Support tens of bytes of packet metadata

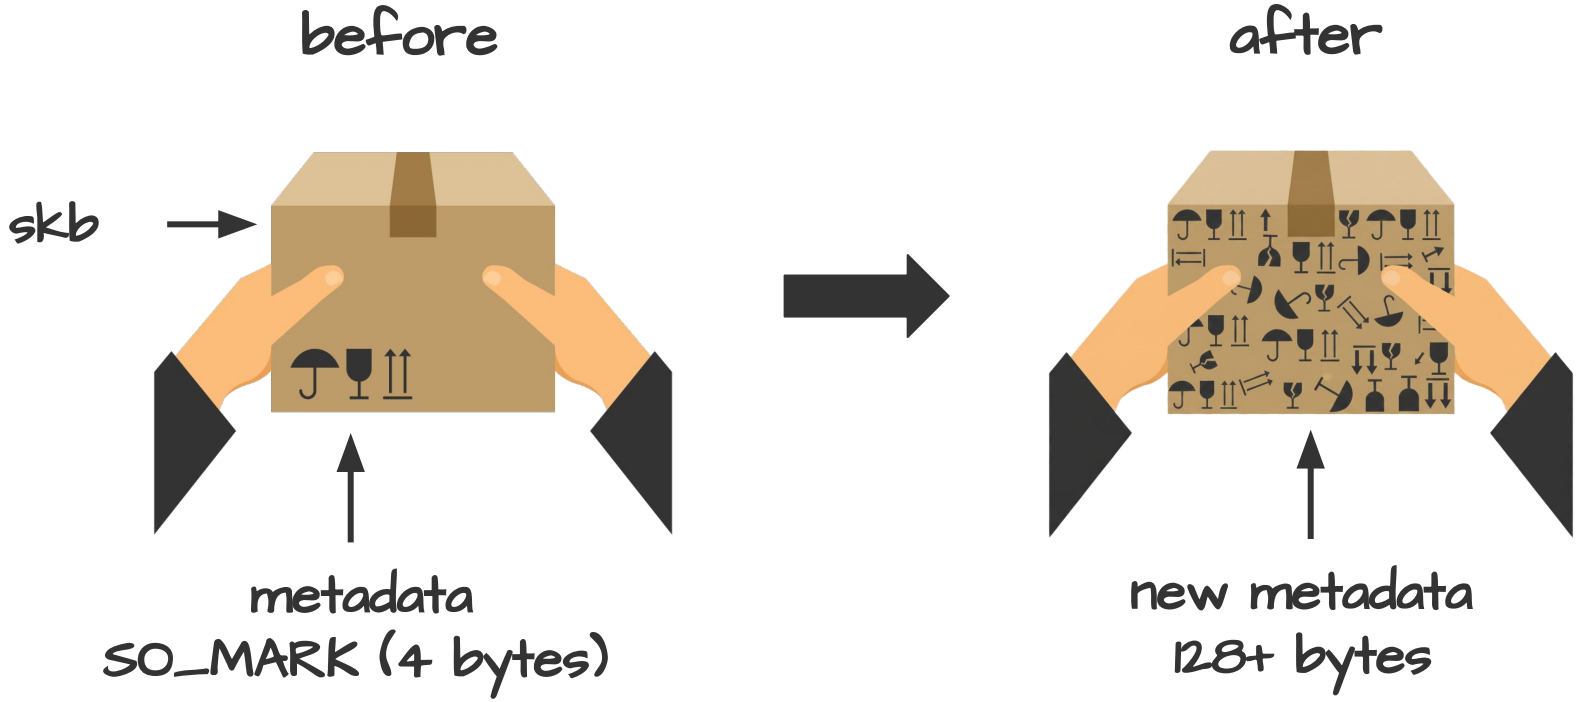

Motivation – Use case #1
Configuration everywhere – Inconsistency everywhere

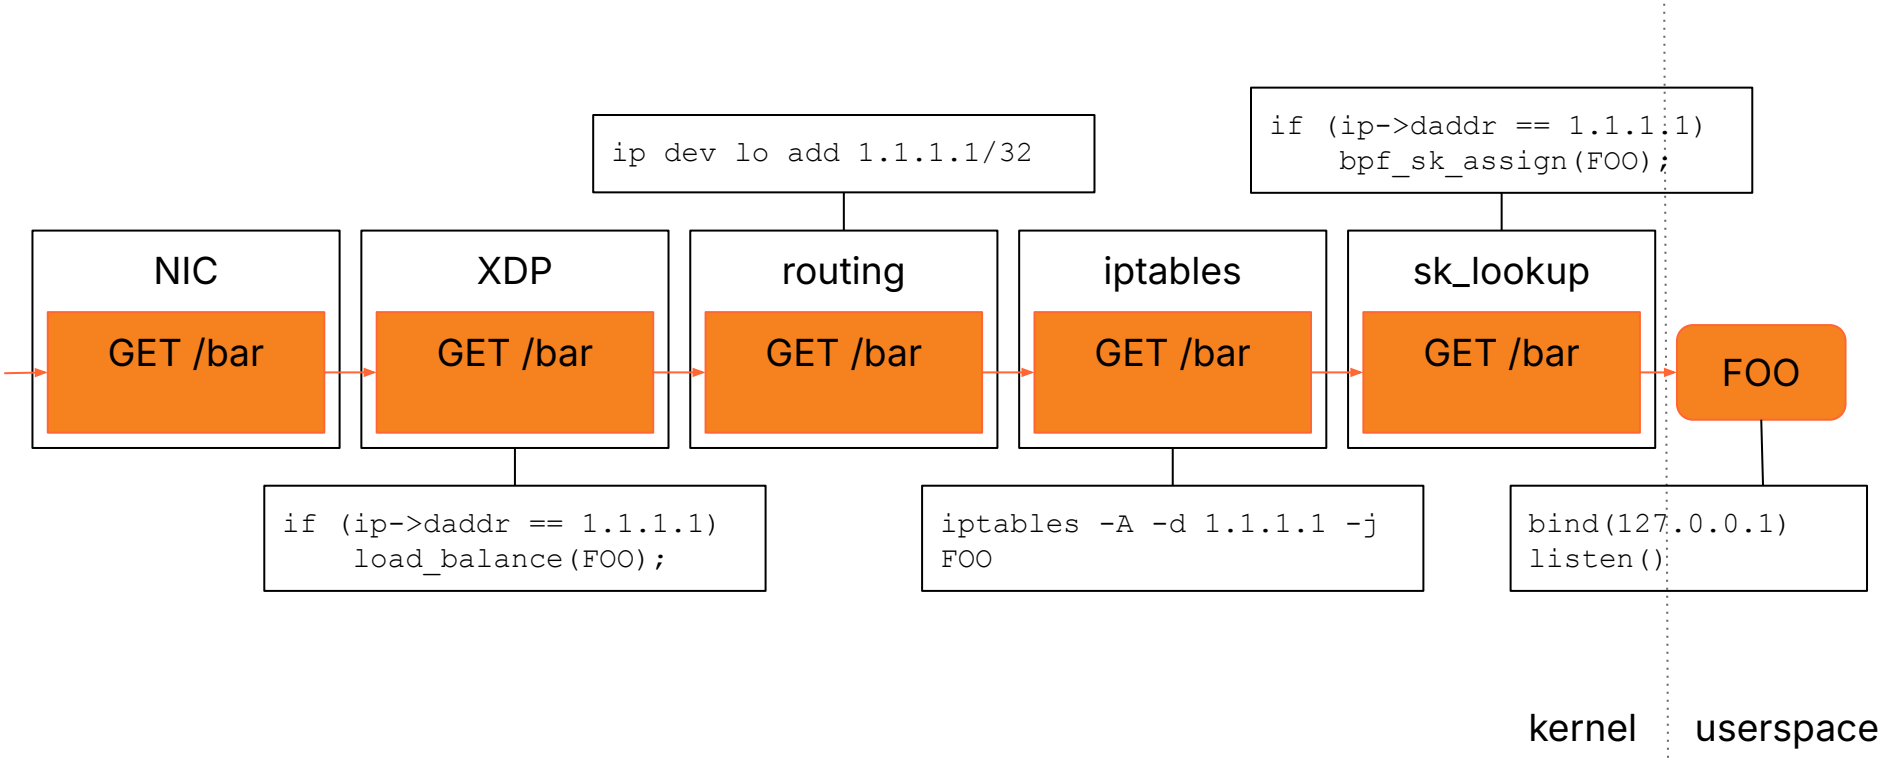

Motivation – Use case #1 – Classify packets in XDP

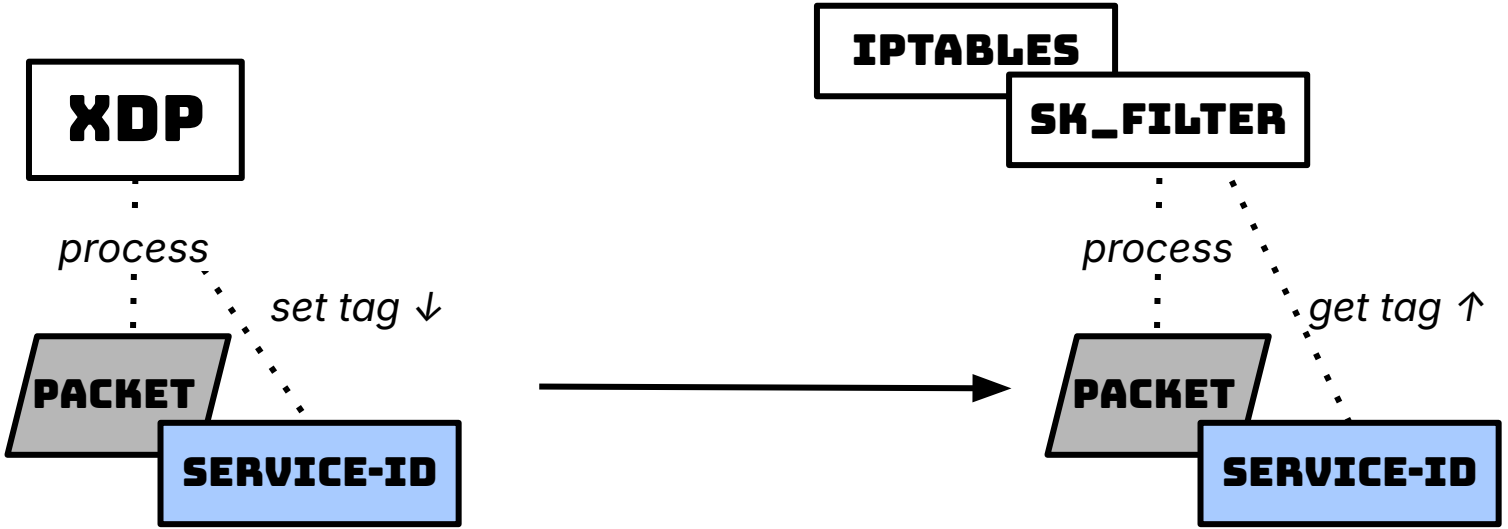

Motivation – Use case #2
Where did this request come from?

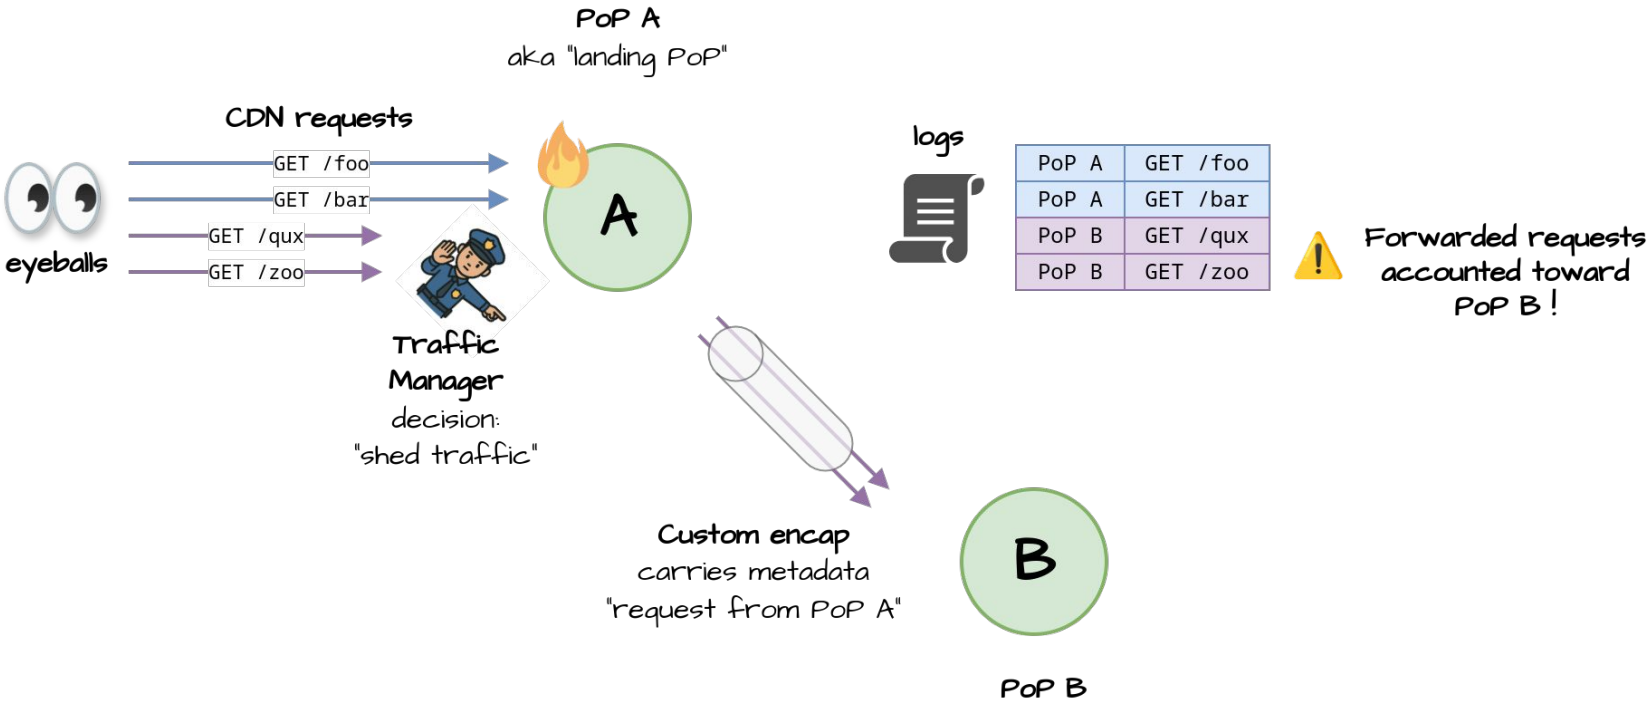

Motivation – Use case #2 – Pass info from encap headers to user-space

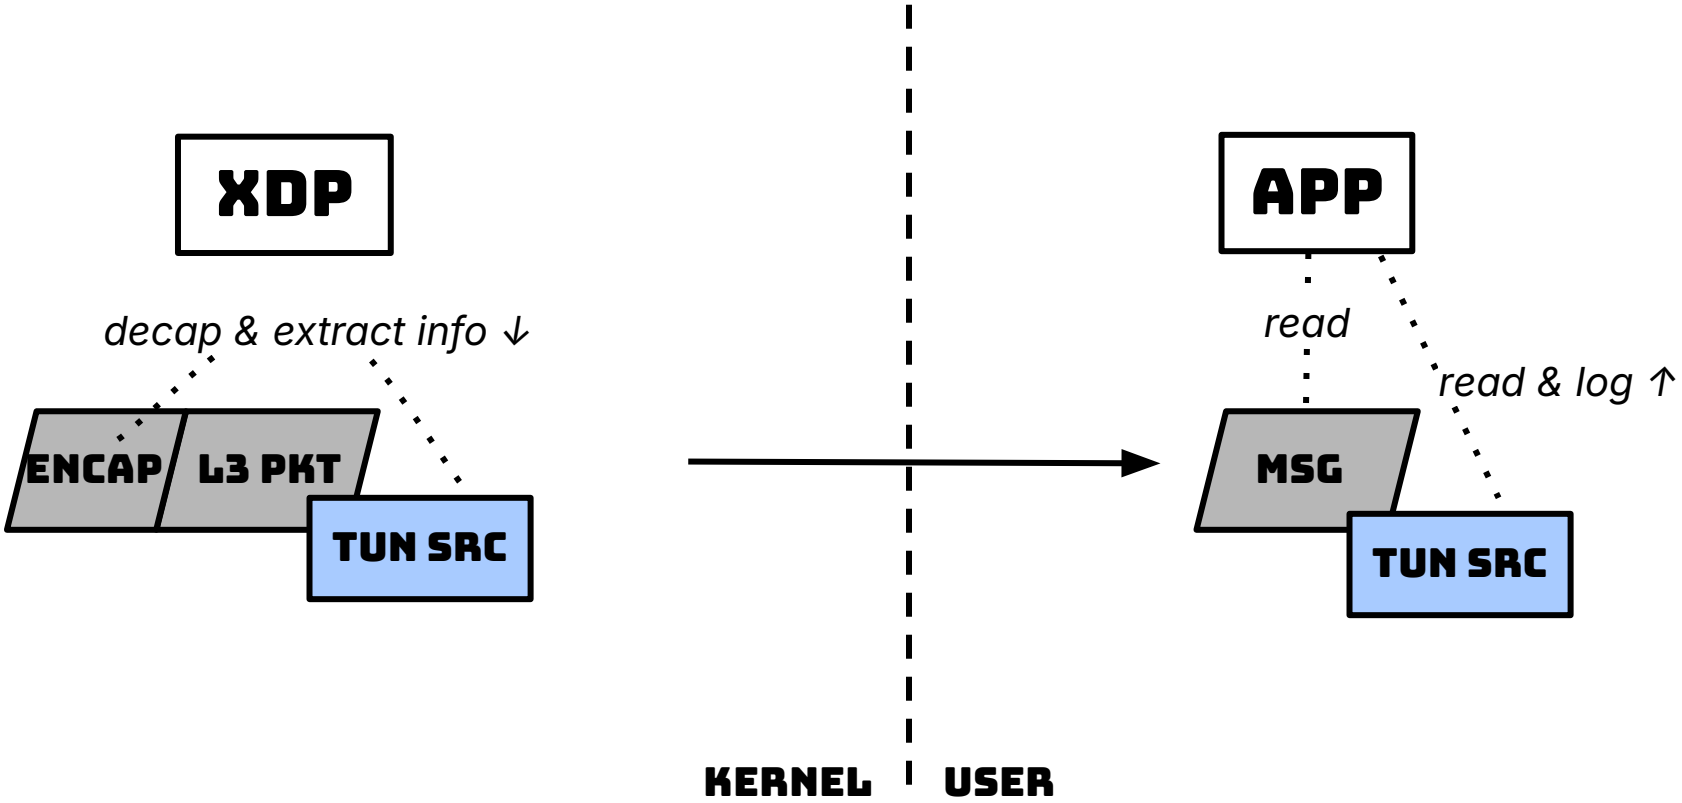

Motivation – Use case #3 – Attribute packet drops

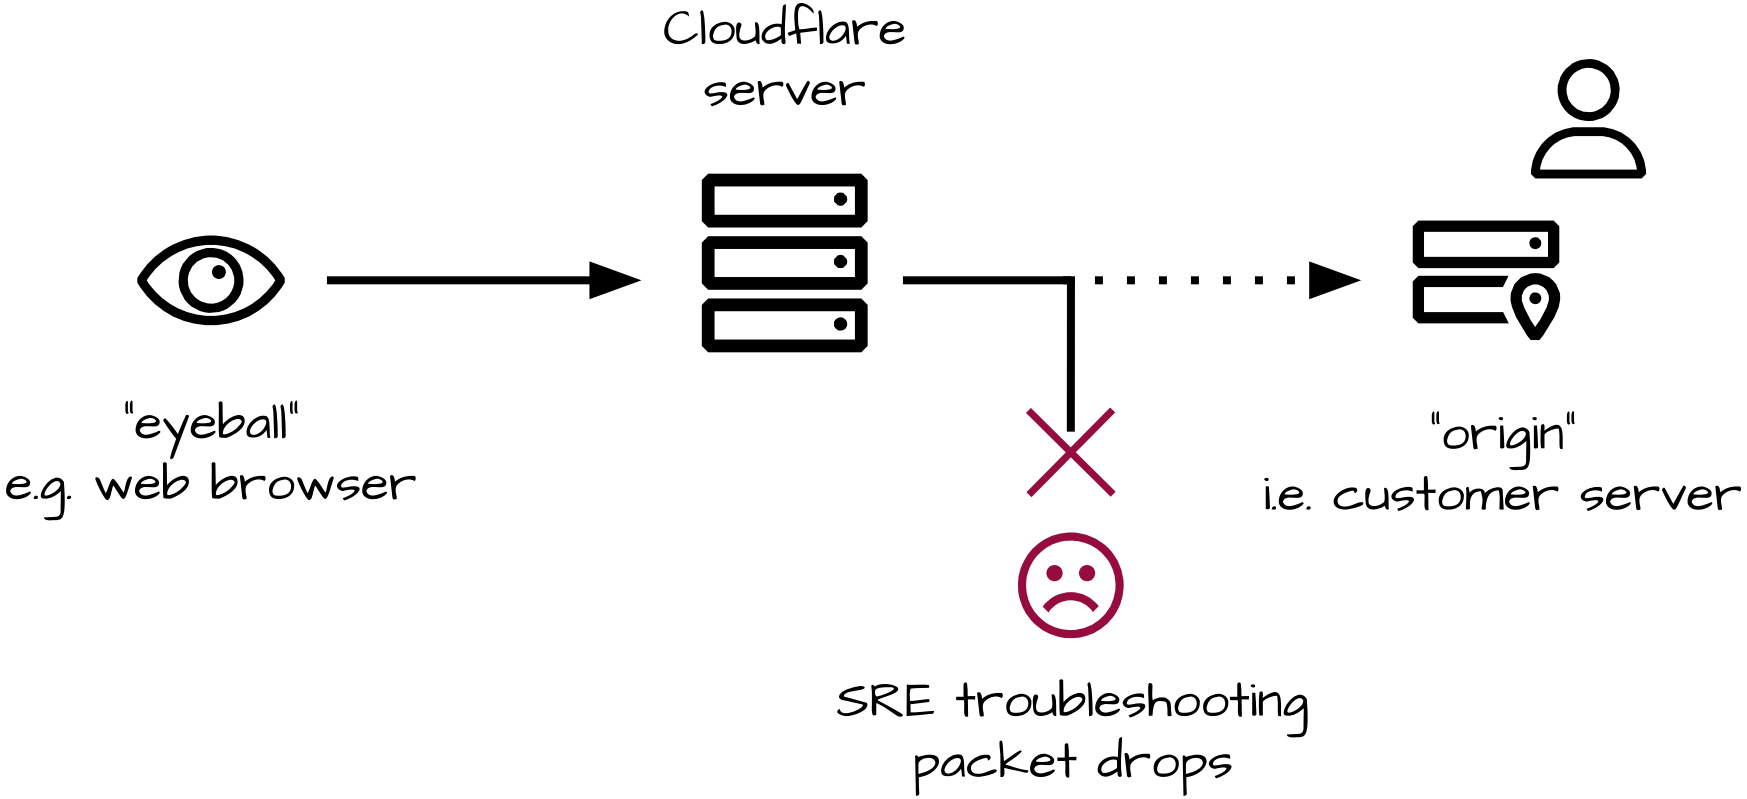

Motivation – Use case #3 – Attribute packet drops

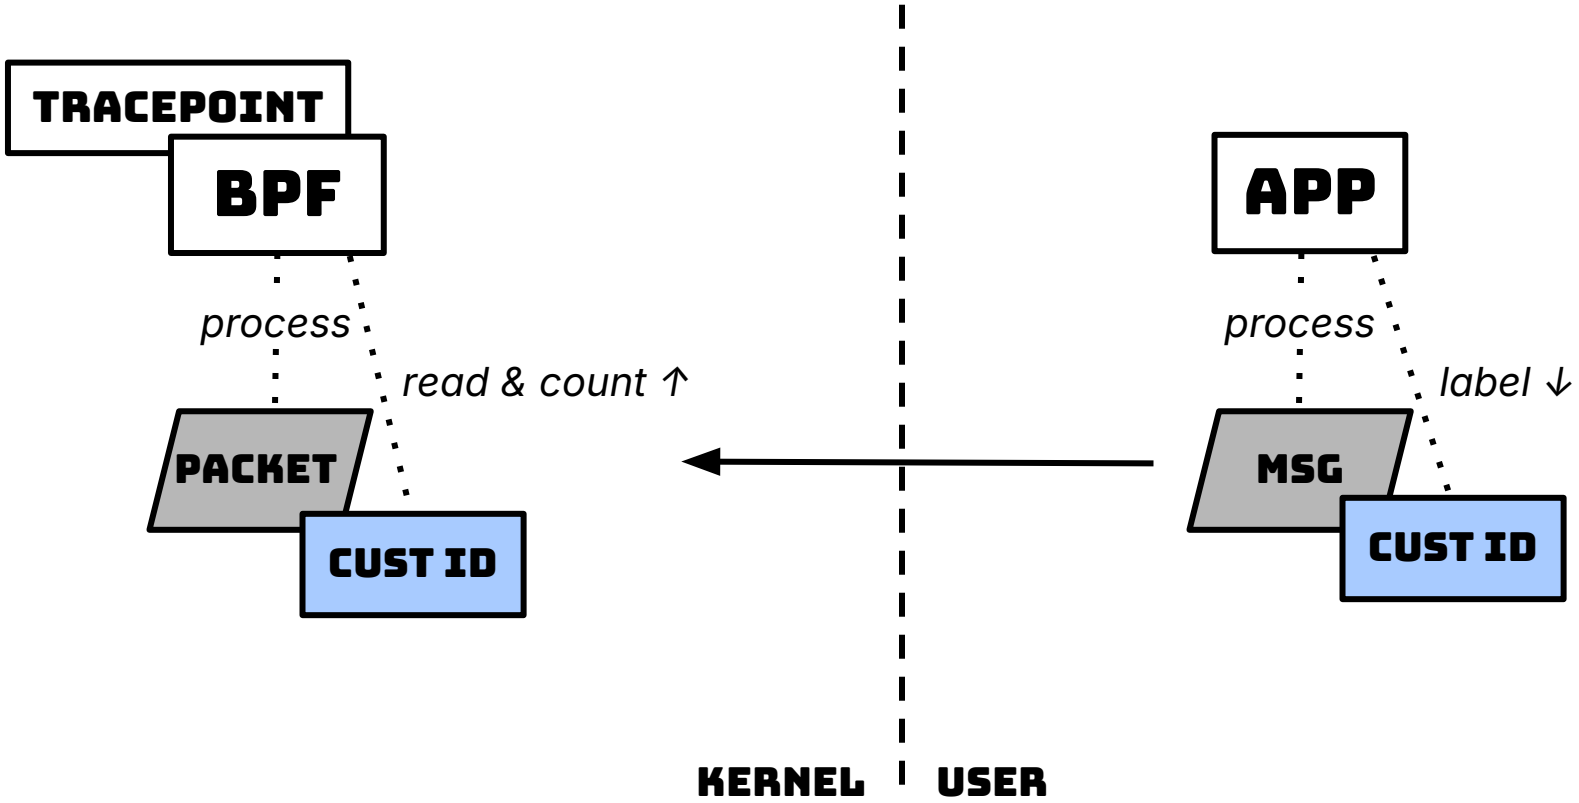

Motivation – Use case #4 – Trace requests

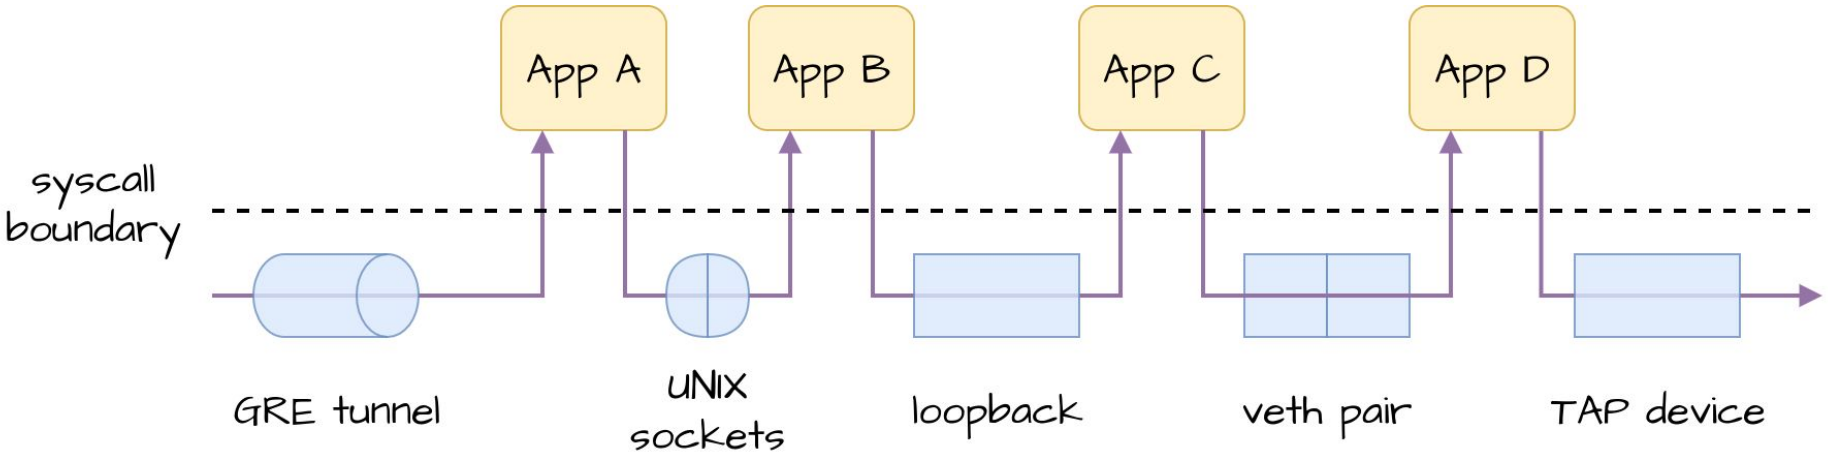

How to attach metadata to packets?


LOST CREEK 55
NOWHERE 10 km
SKB TRAITS 42 km

packet buffer when XDP attached

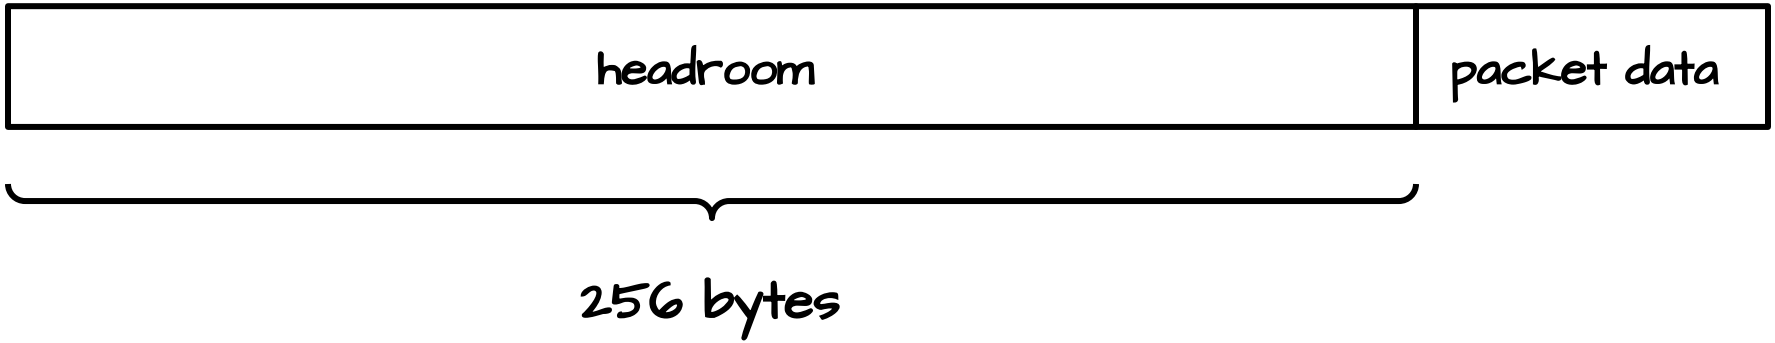

packet buffer

↓

skb head buffer

skb->head

skb->data


xdp_frame{}
headroom
packet data


xdp_frame{}
headroom
XDP metadata
packet data

- inaccessible past TC BPF hook
- gets corrupted when pushing headers

- accessible everywhere
- is not in the way of headers

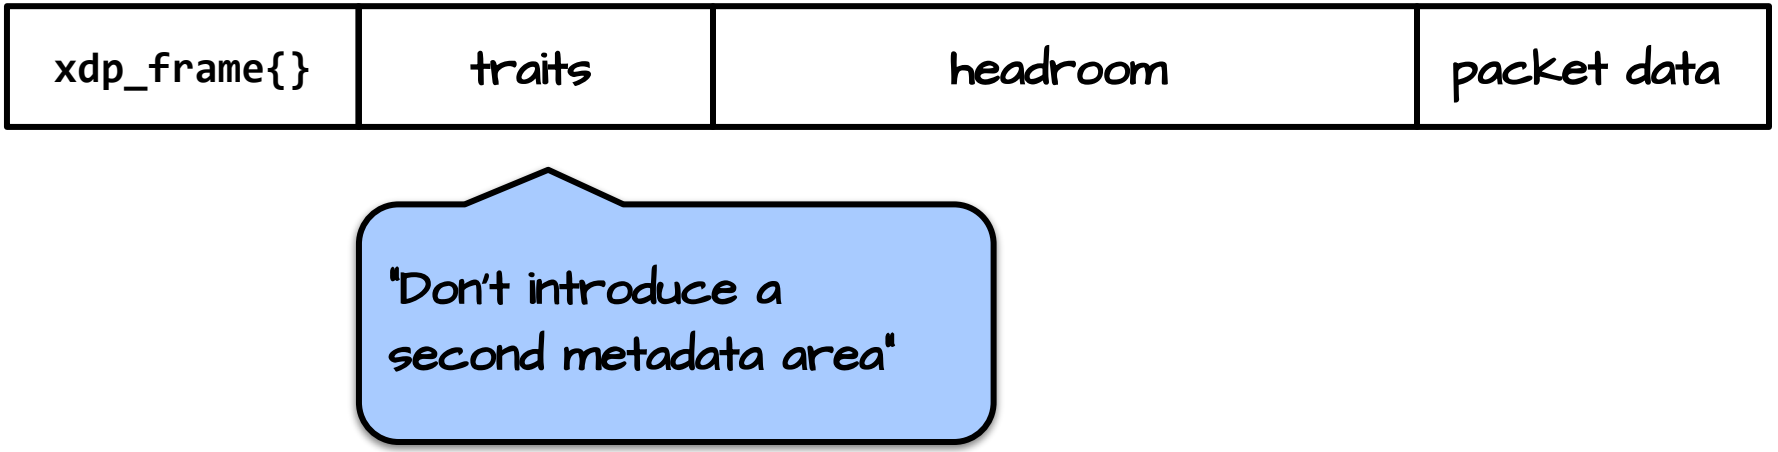
xdp_frame{}
traits
headroom
packet data
"Don't introduce a second metadata area"

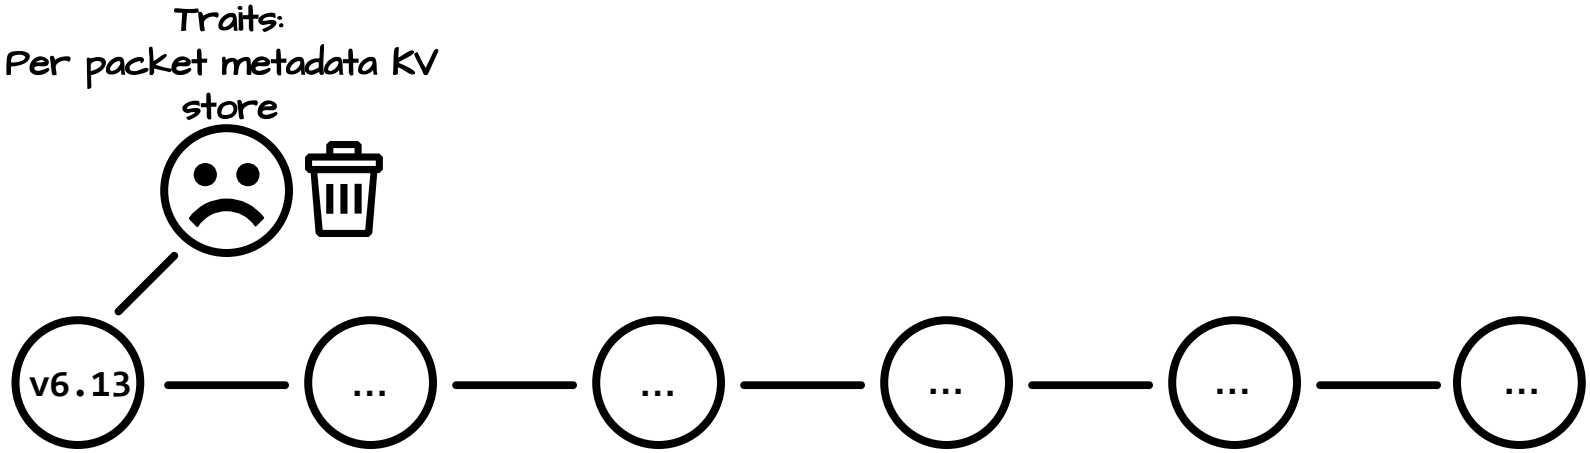
Traits:
Per packet metadata KV store
v6.13
...
...
...
...
...


SKB TRAITS
NOWHERE
XDP METADATA

CLOUDFLARE

How do we fix XDP/skb metadata?*

* and stay backwards compatible

What do we need to make `skb->data_meta` work?

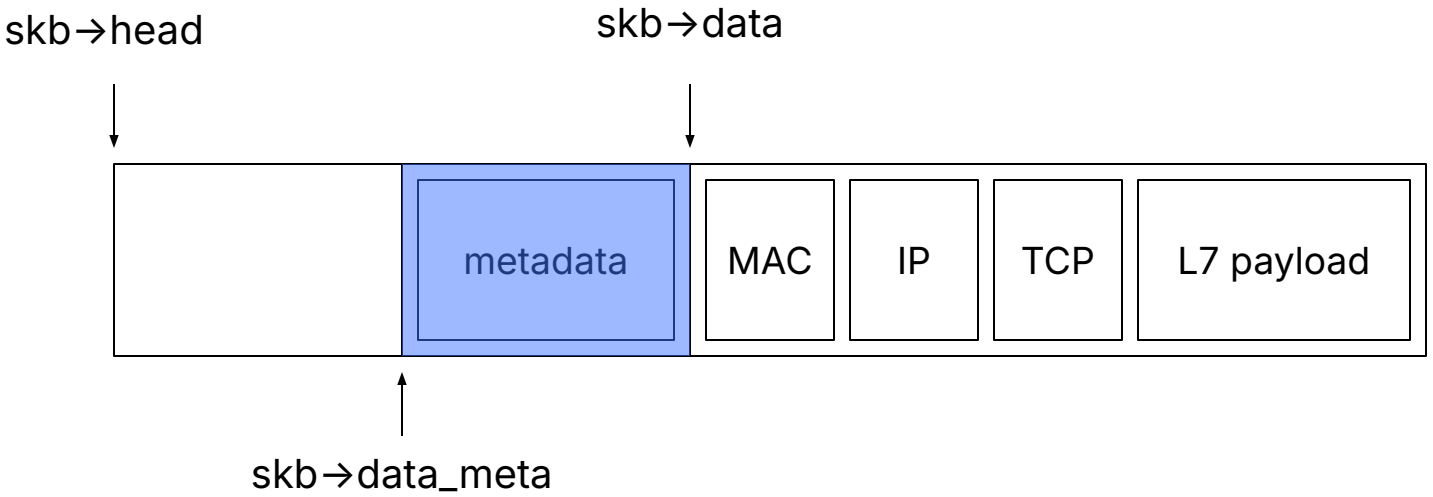

What do we need to make `skb->data_meta` work?

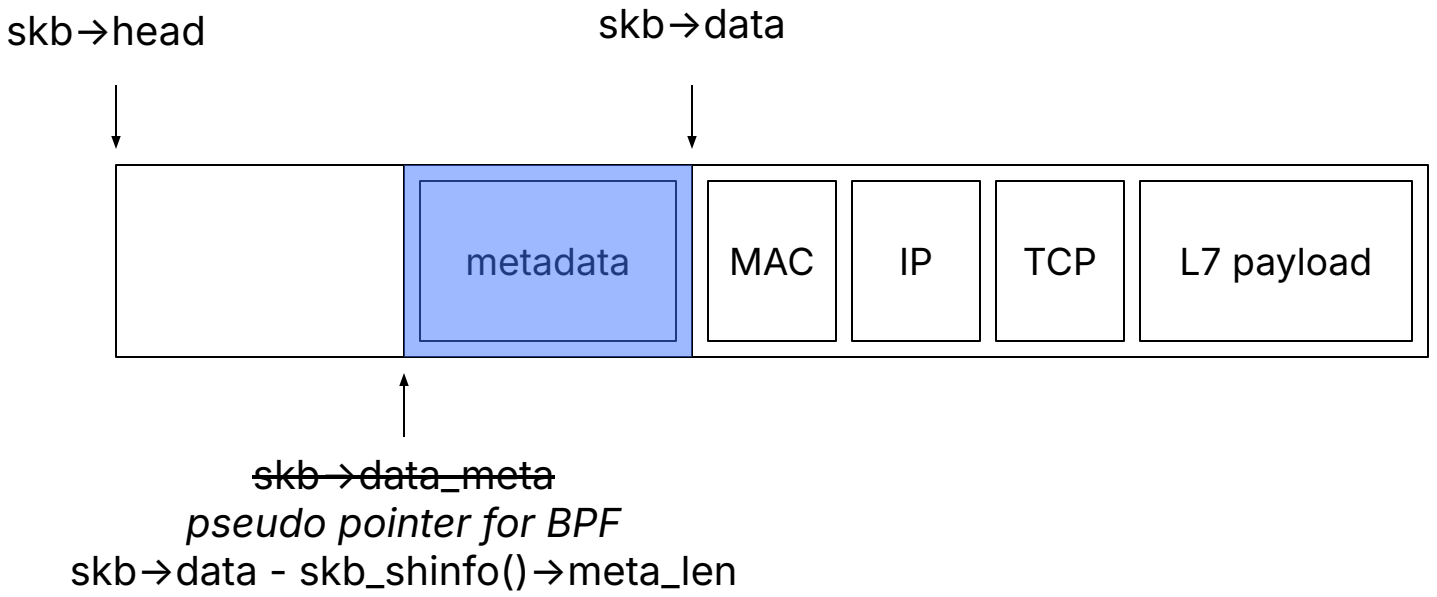

As we pull packet headers...

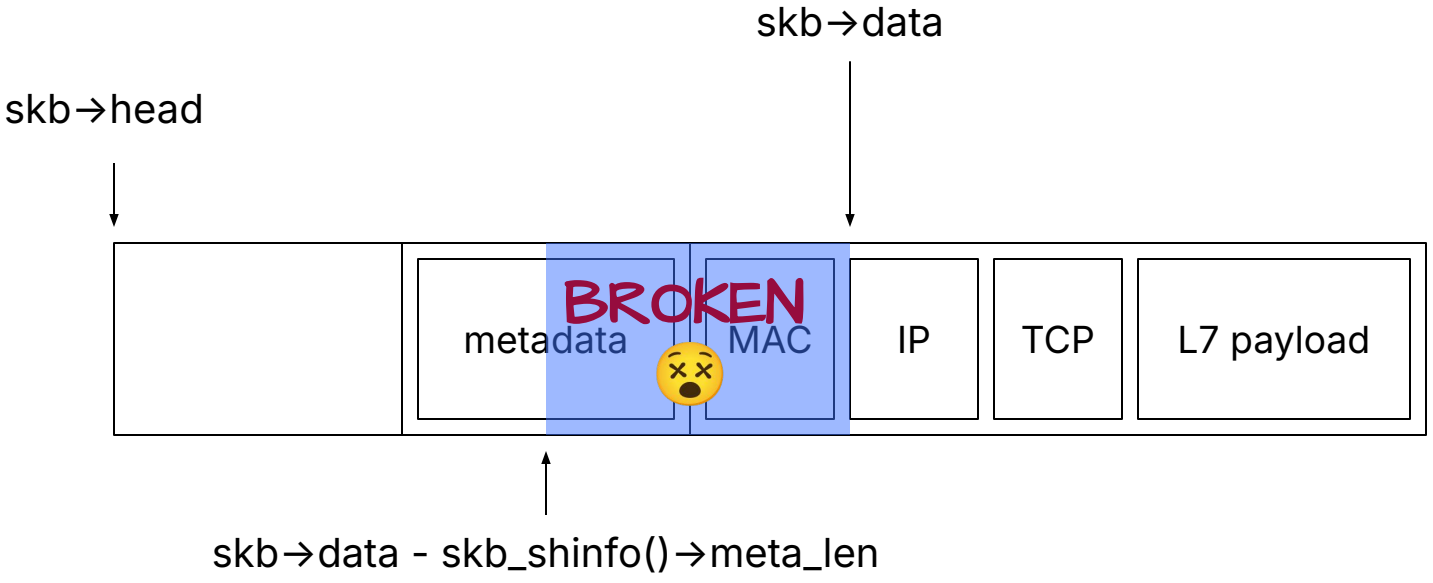

… we'd need to keep moving the metadata

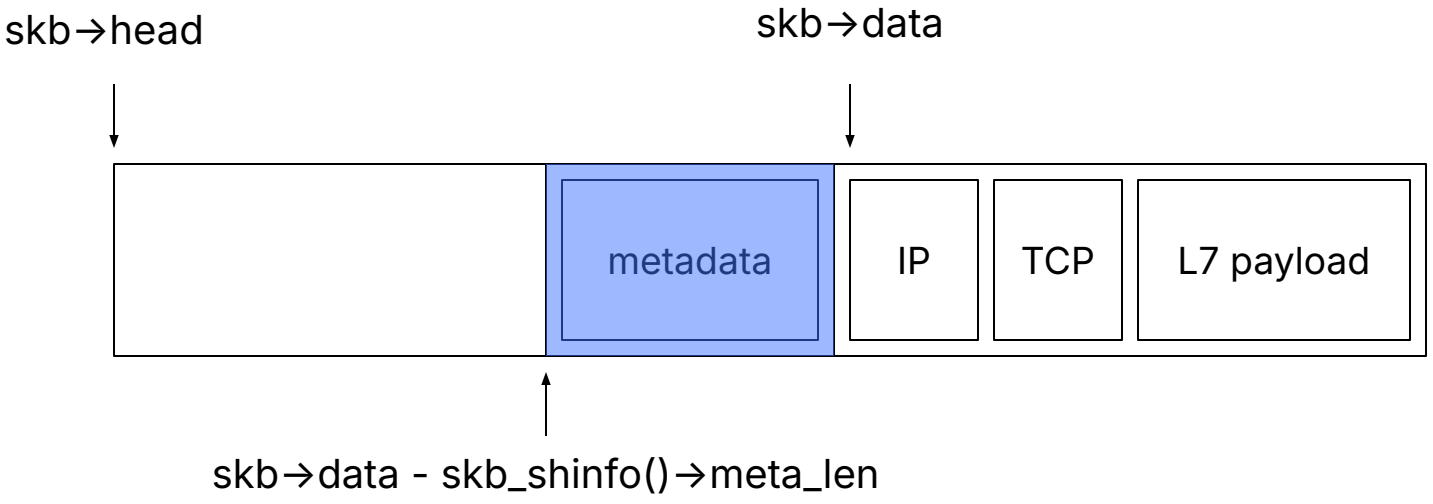

Alternative?
Another layer of indirection

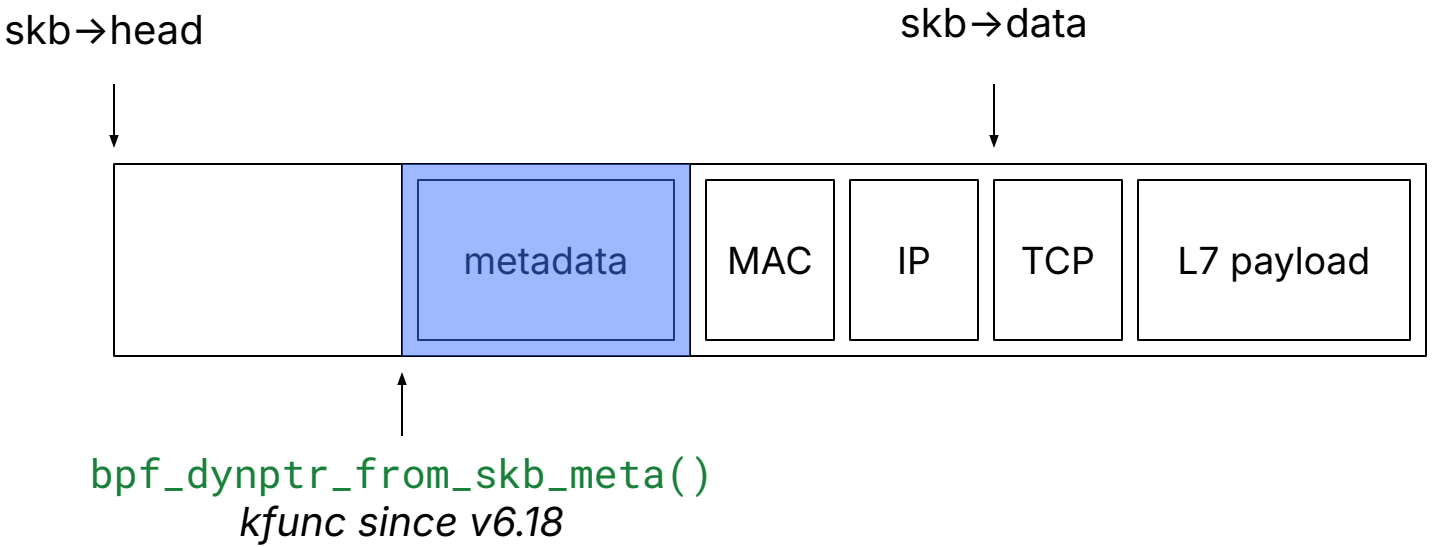

https://msgid.link/20250814-skb-metadata-thru-dynptr-v7-0-8a39e636e0fb@cloudflare.com/

It does not rely on `skb->data` ...

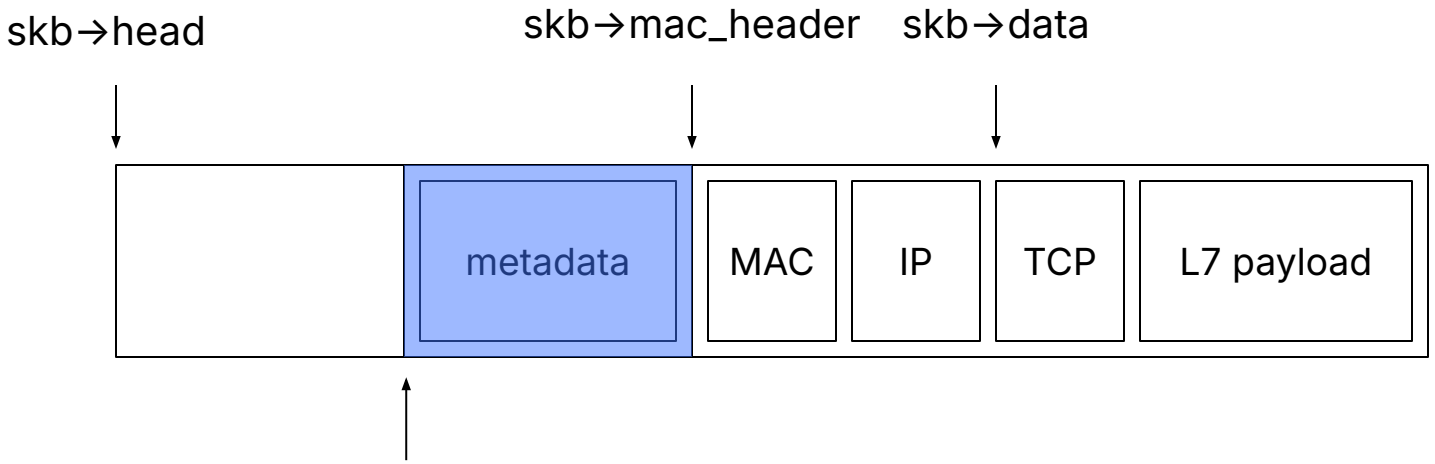

Add a dynptr type for skb metadata for TC BPF

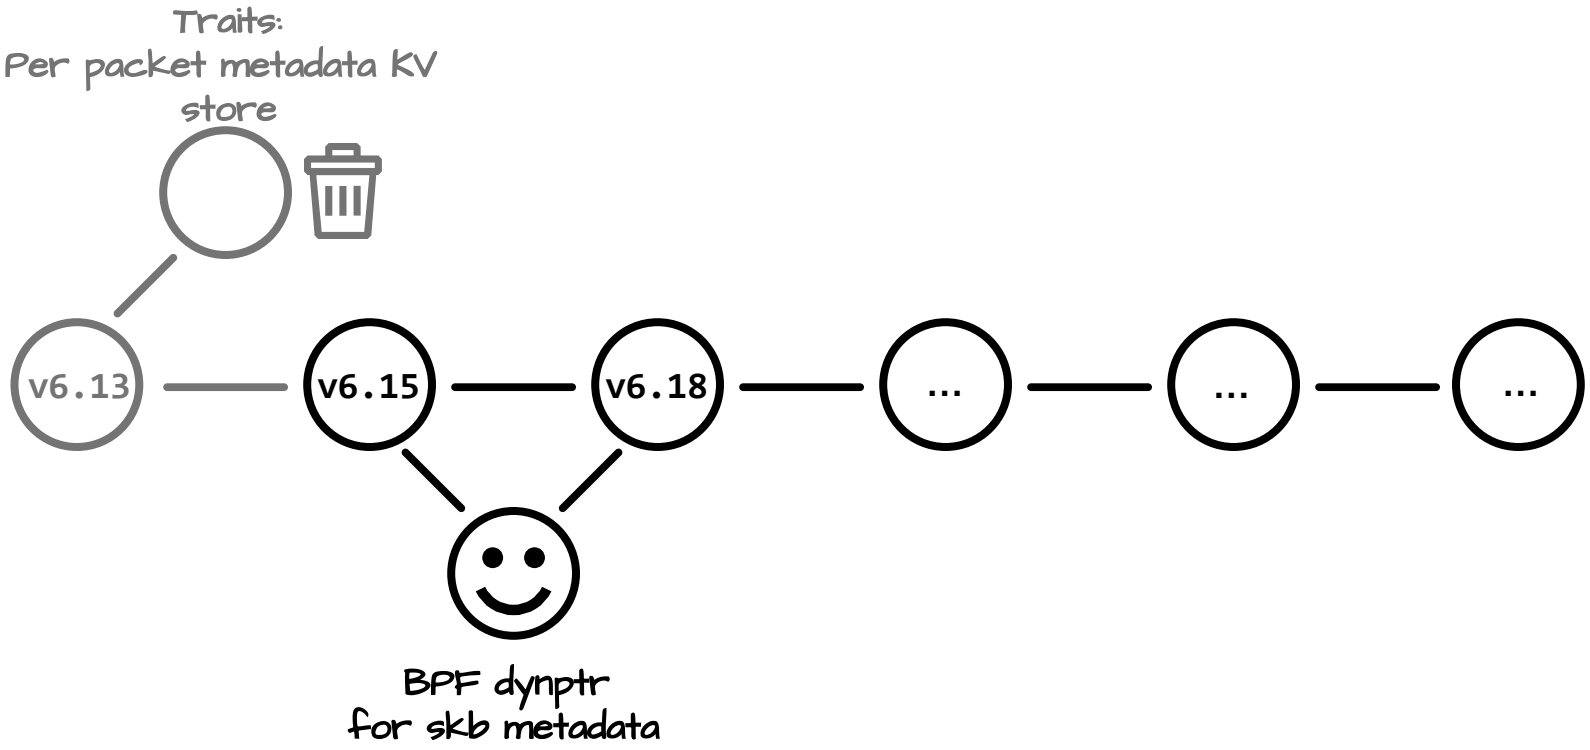

Make TC BPF helpers preserve skb metadata

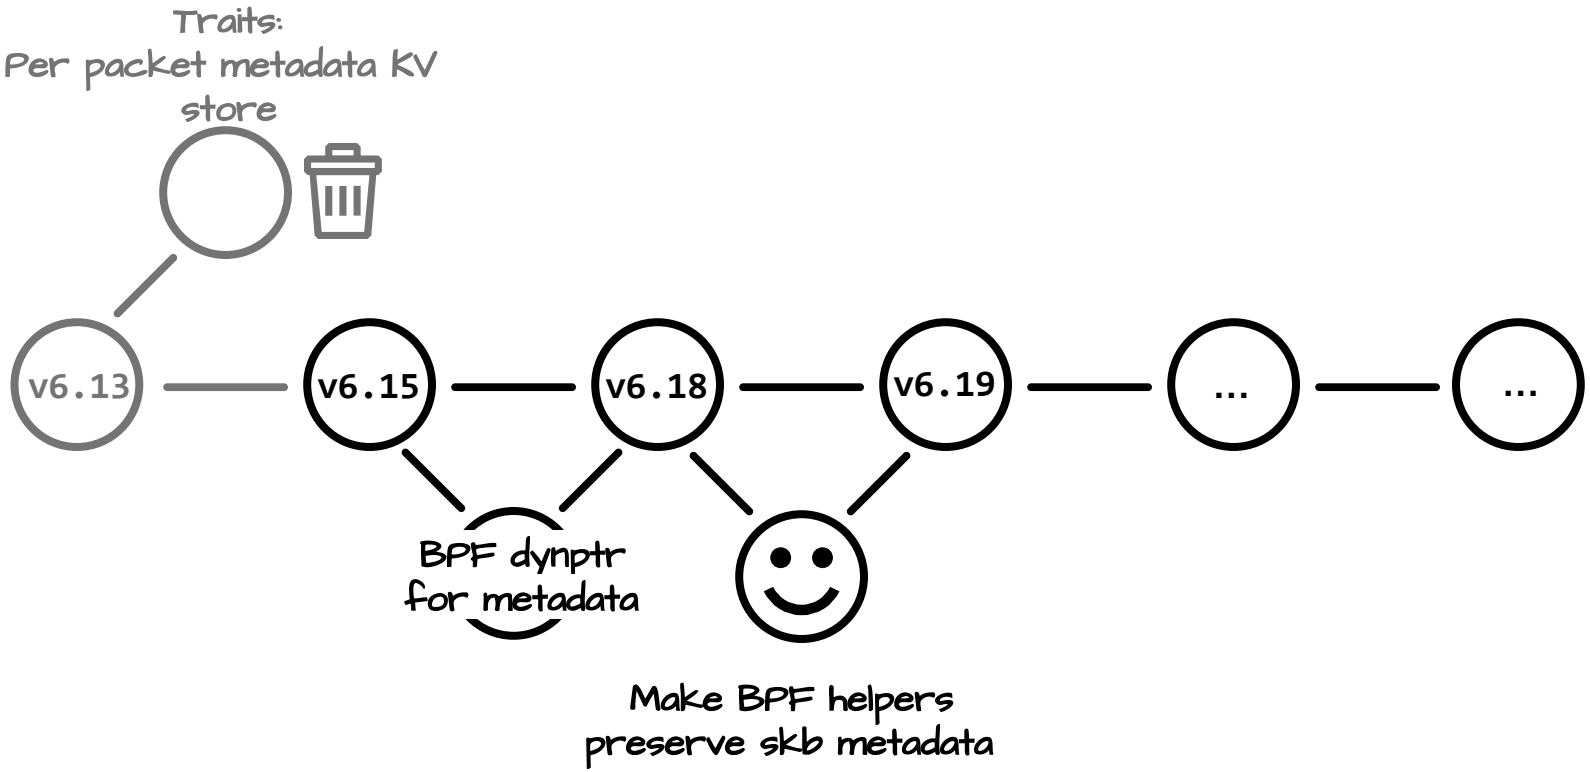

But if we have L2 tunnels...

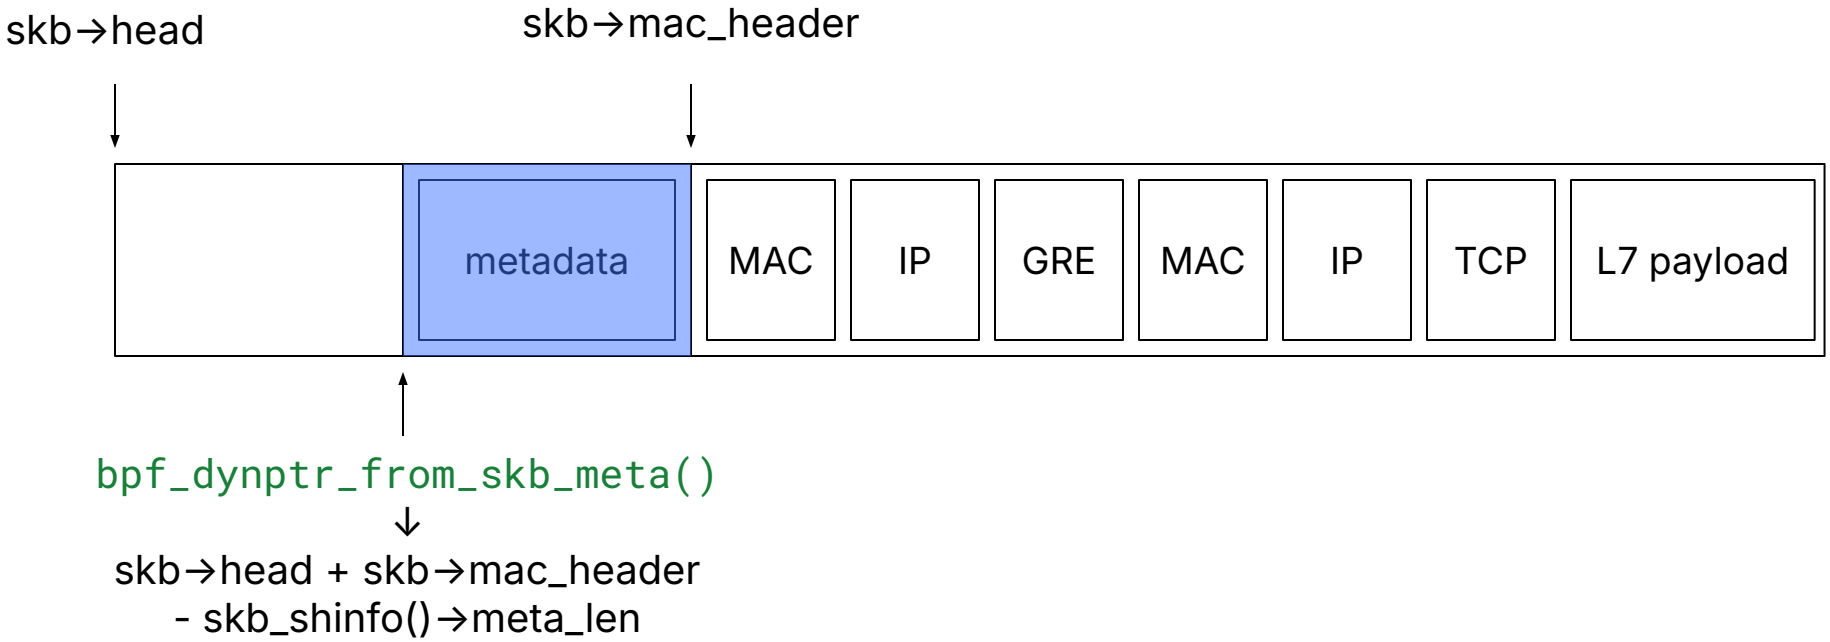

… it breaks down when you reset the MAC header

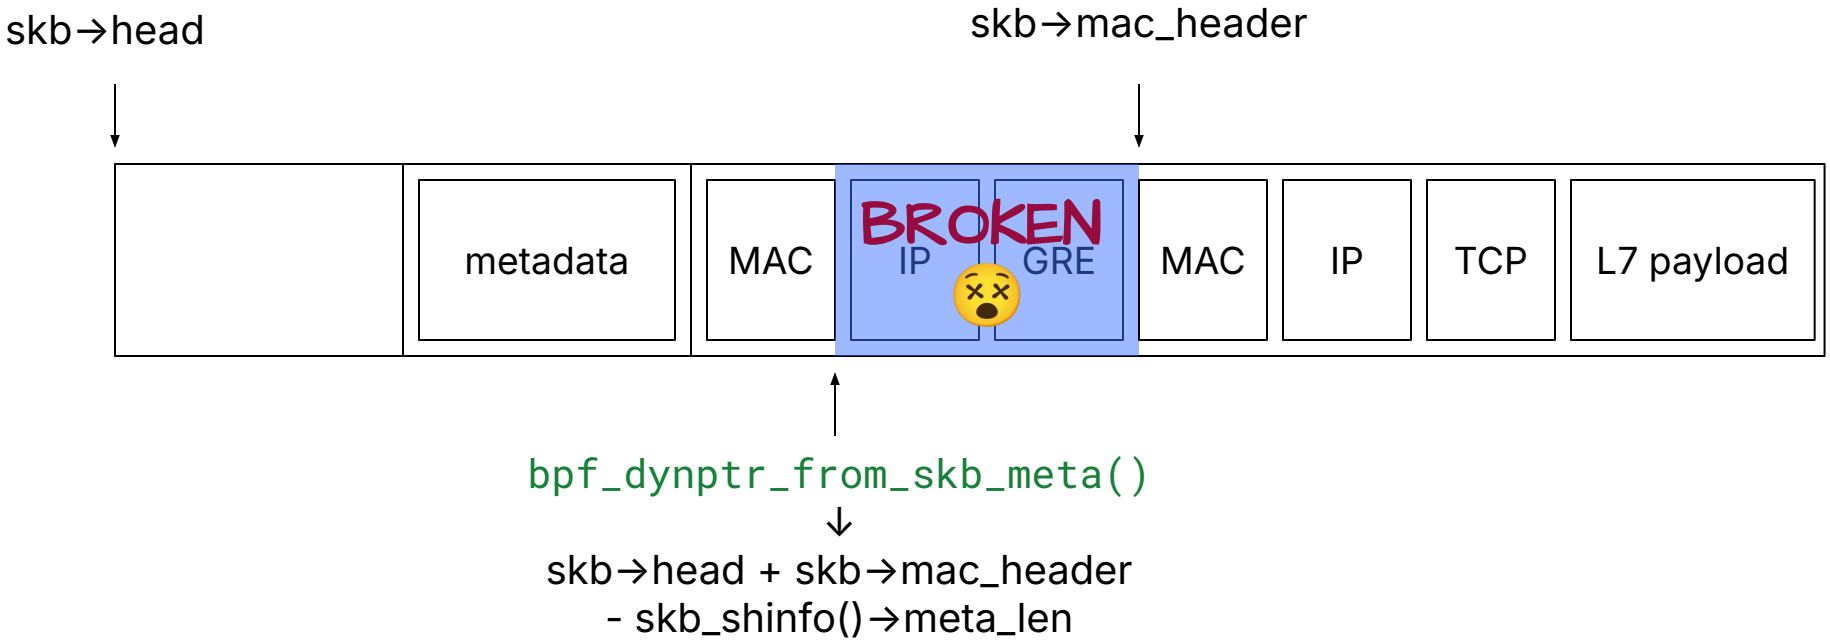

Solution?
Track metadata offset separately from MAC

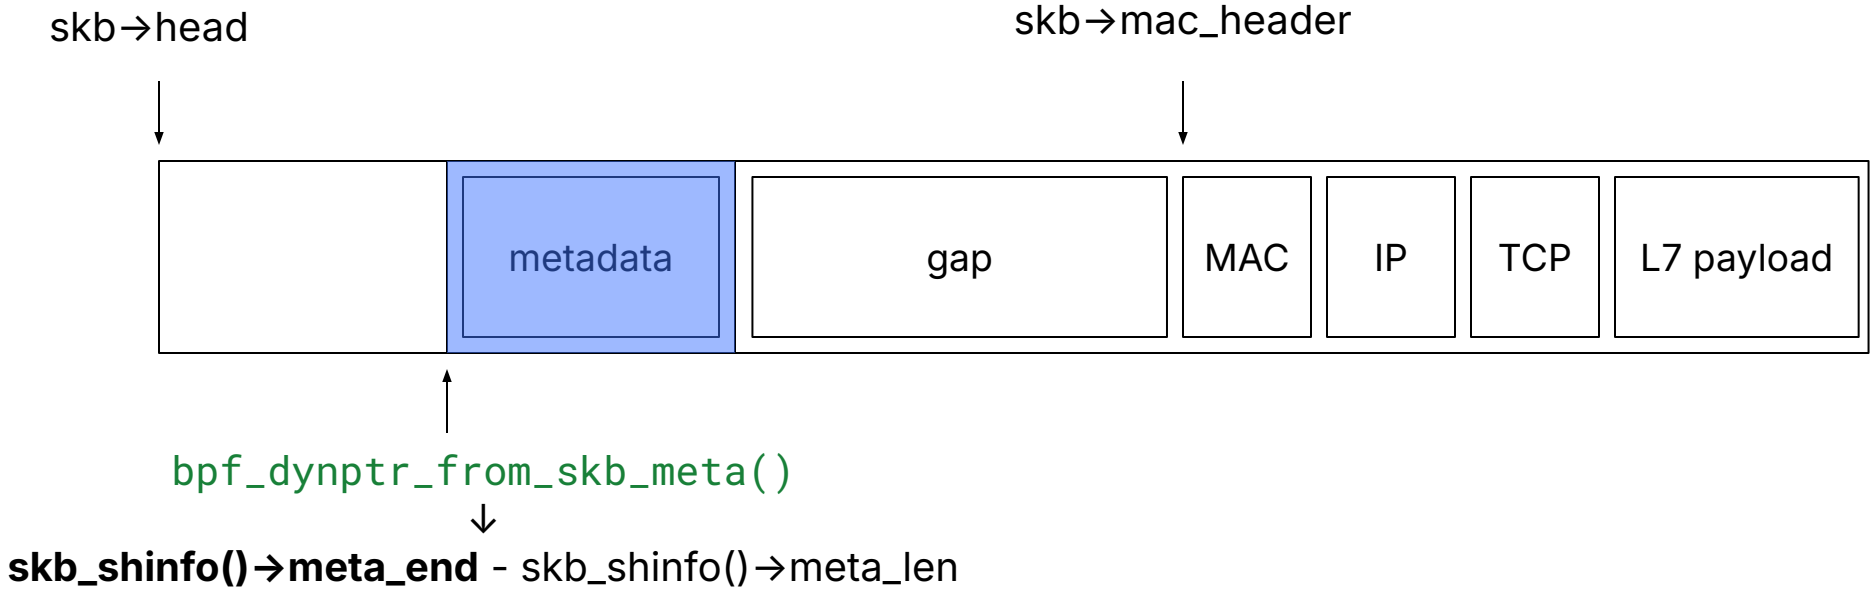

Can move the metadata out of the way when pushing headers (TX path)

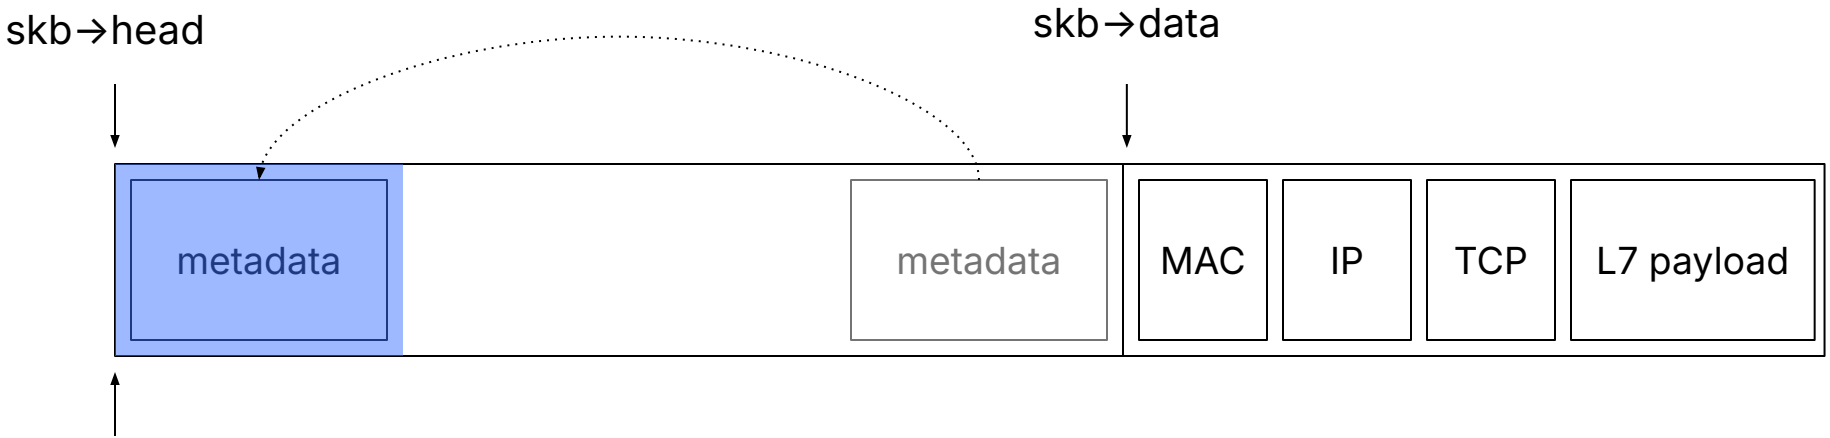

`bpf_dynptr_from_skb_meta()`

↓

skb_shinfo()→meta_end - skb_shinfo()→meta_len

The curse of `skb->data_meta`

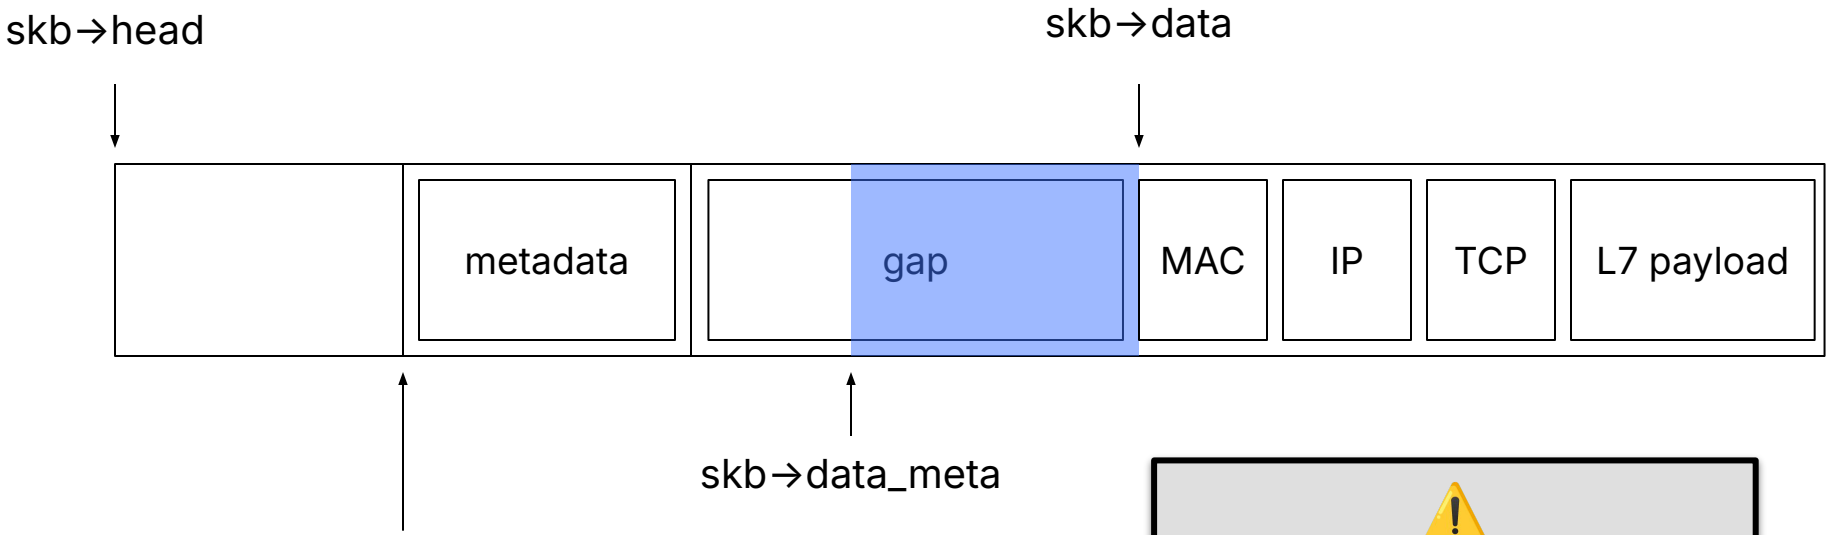

⚠️

We need to keep moving the metadata when using `skb->data_meta` in TC BPF

Proposed layout –
Metadata anywhere within the headroom

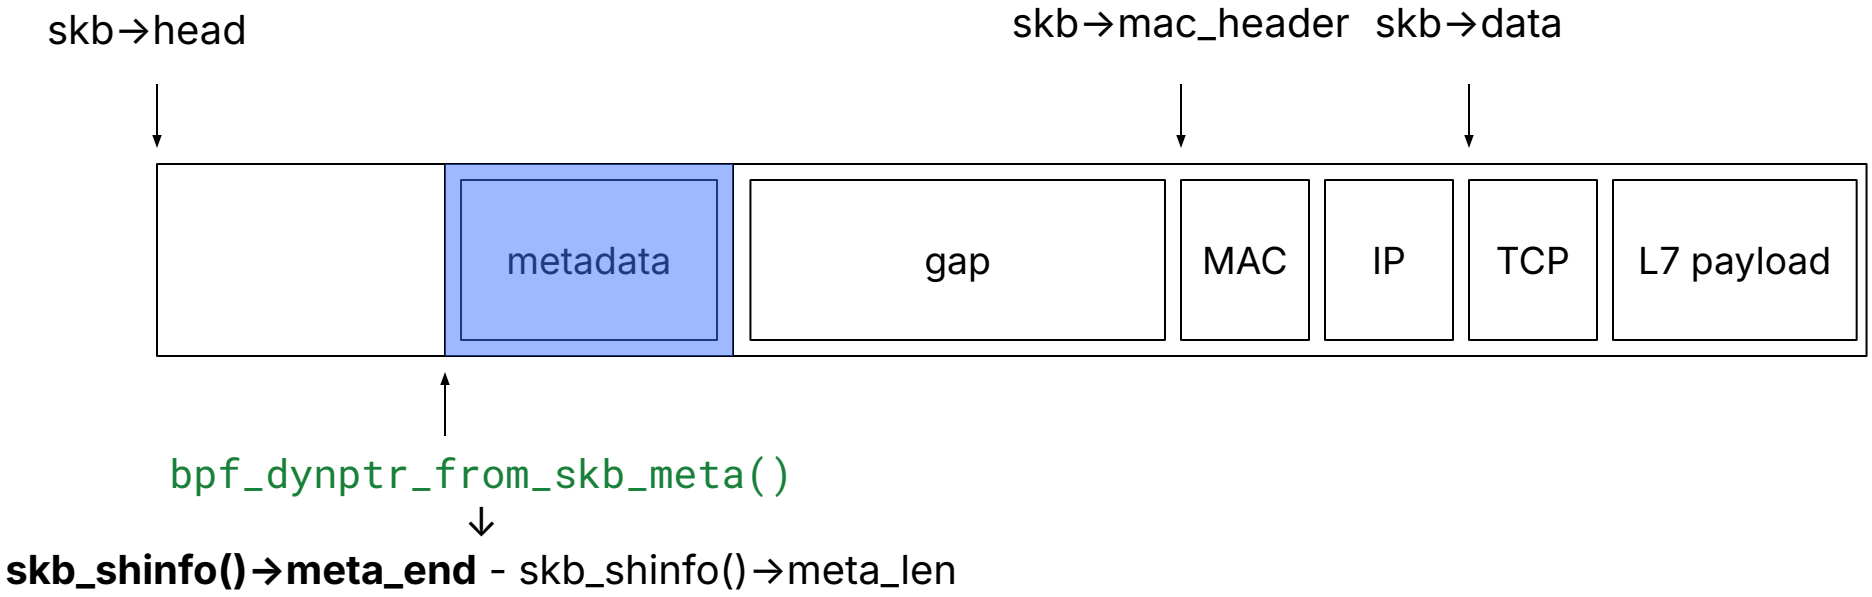

Proposed layout –
Metadata anywhere within the headroom

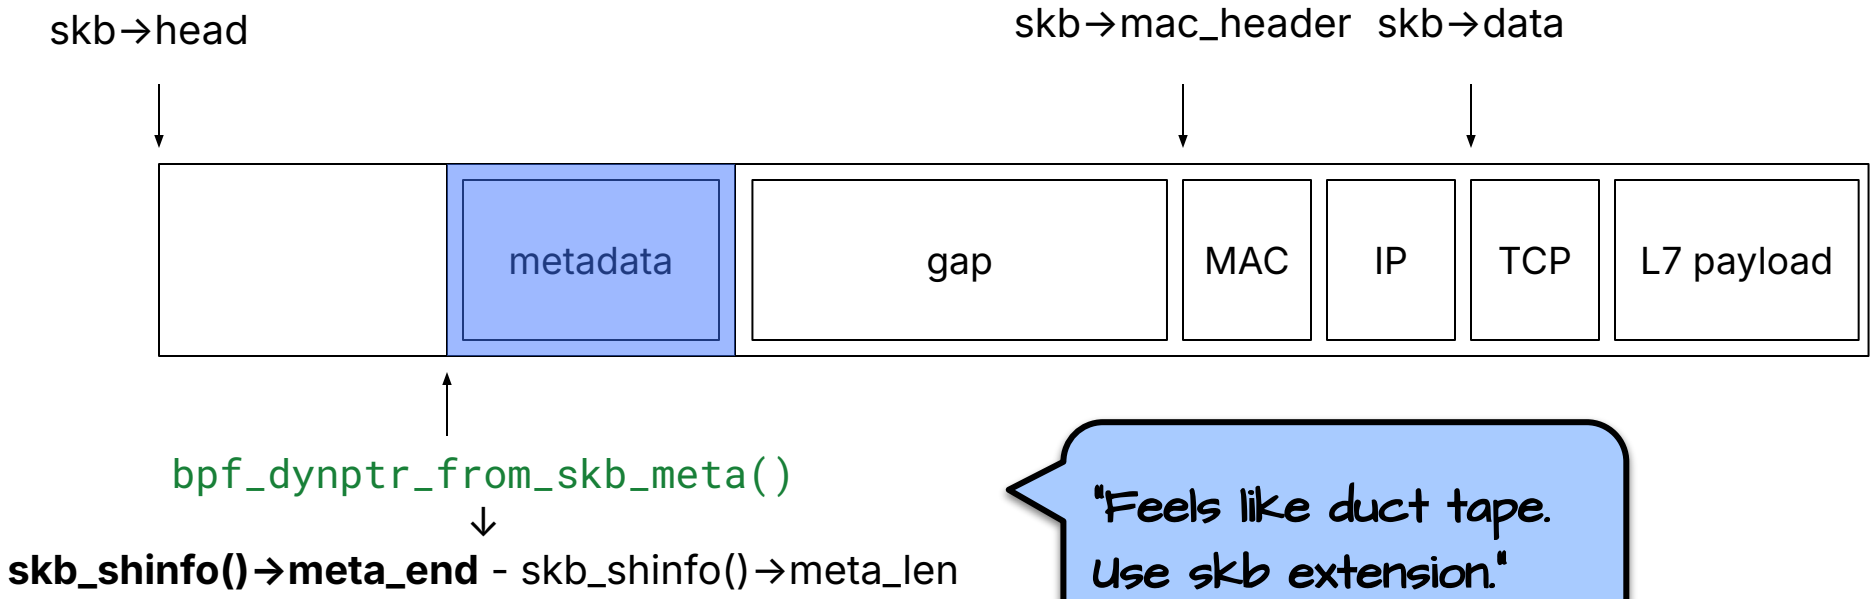

LPC 2025: Decouple skb metadata tracking from MAC header offset

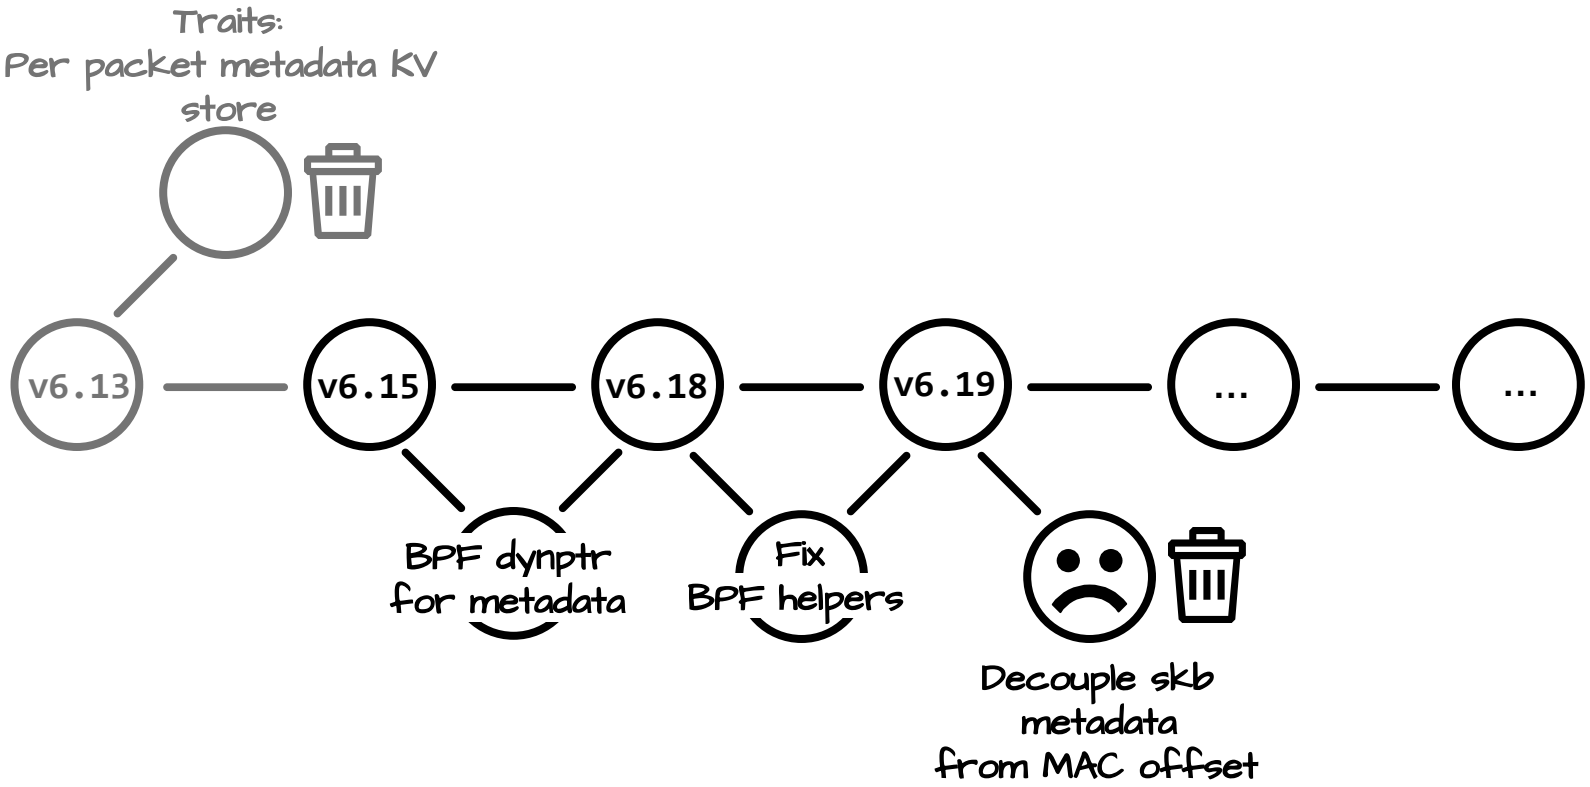


SKB TRAITS
NOWHERE
SKB EXTENSION
XDP METADATA

skb extension

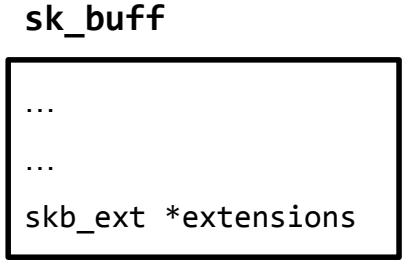

skb extension

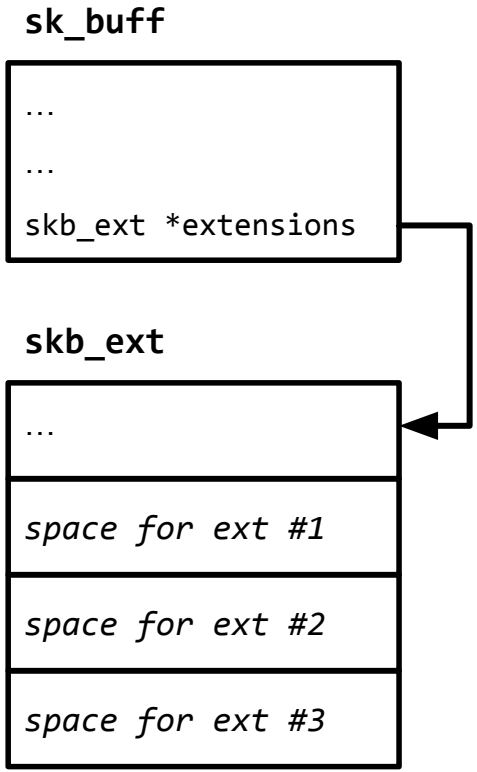

skb extension for what?

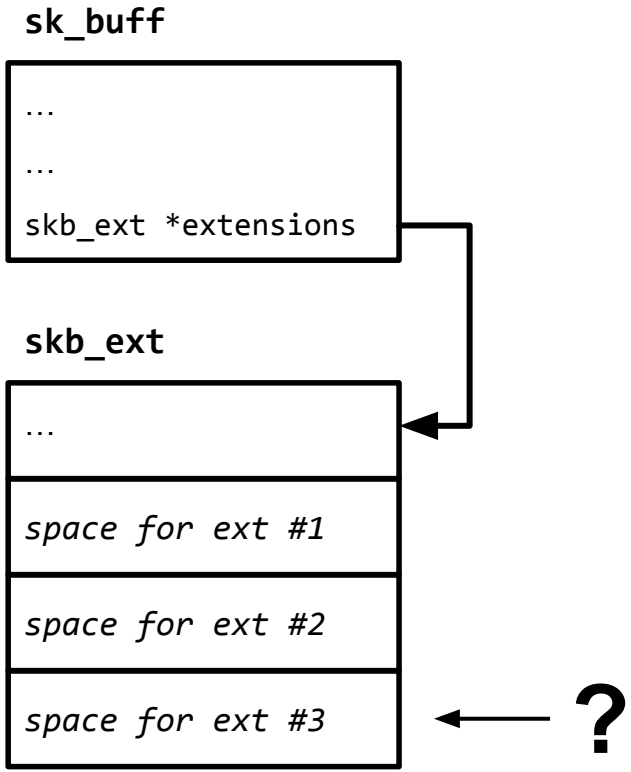

skb extension for a simple metadata buffer?

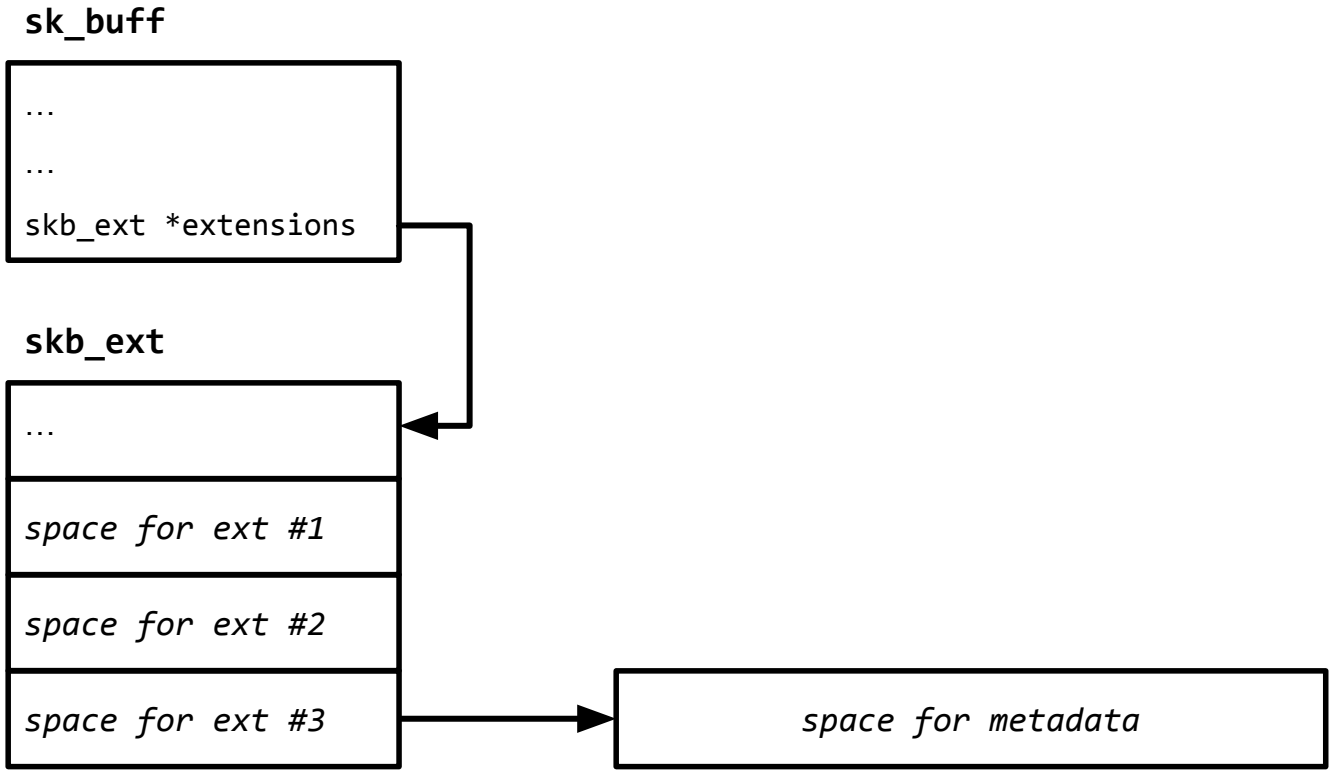

BPF local storage

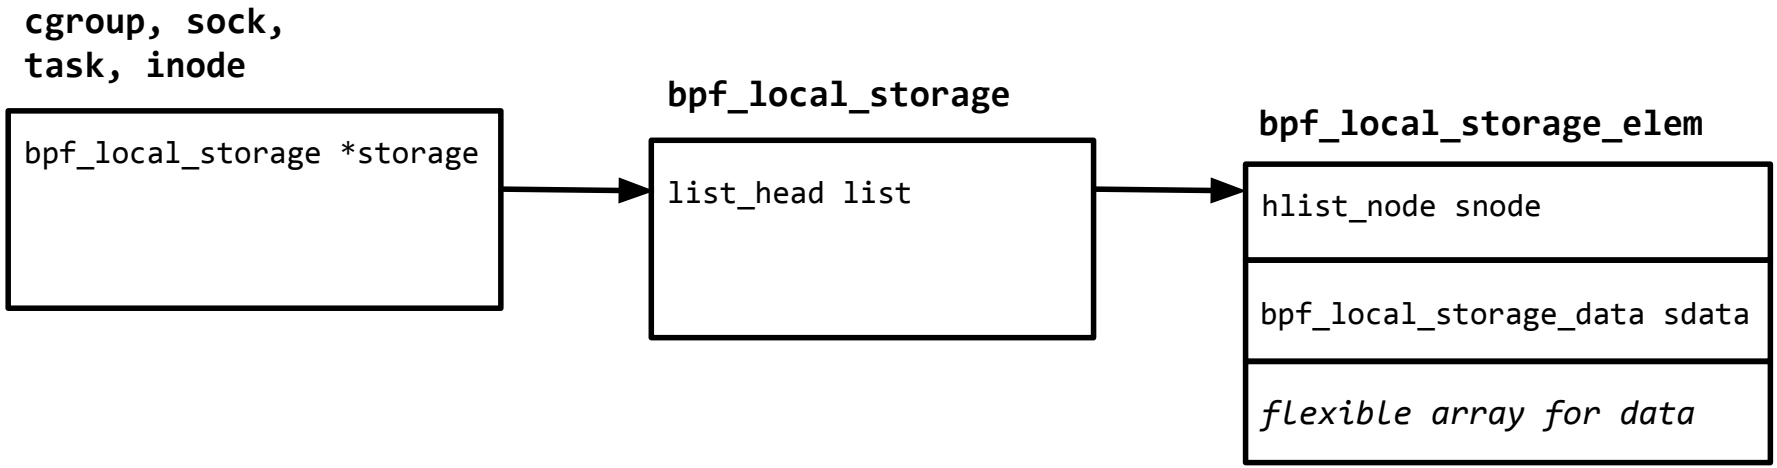

BPF local storage

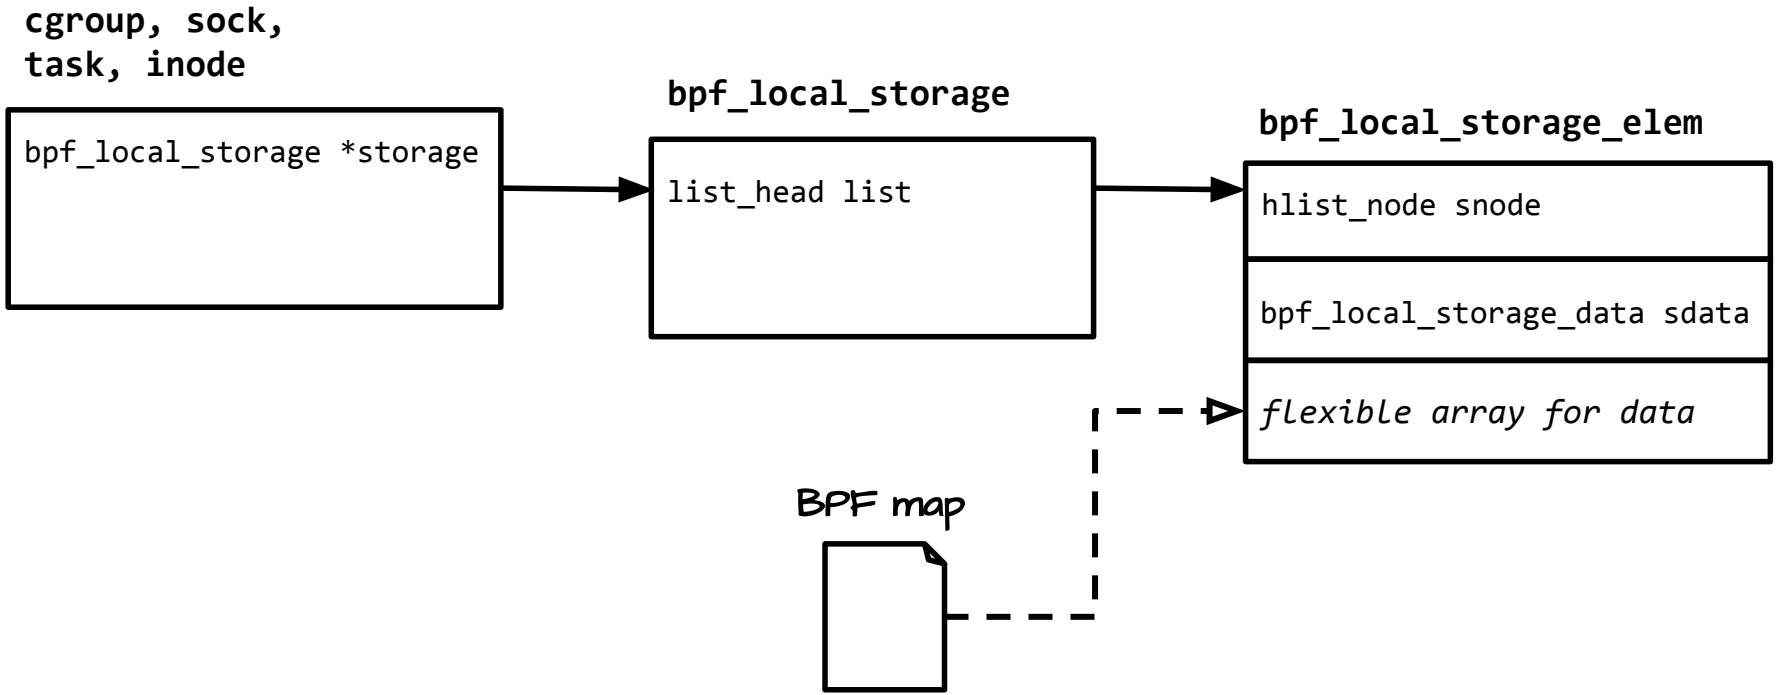

BPF local storage

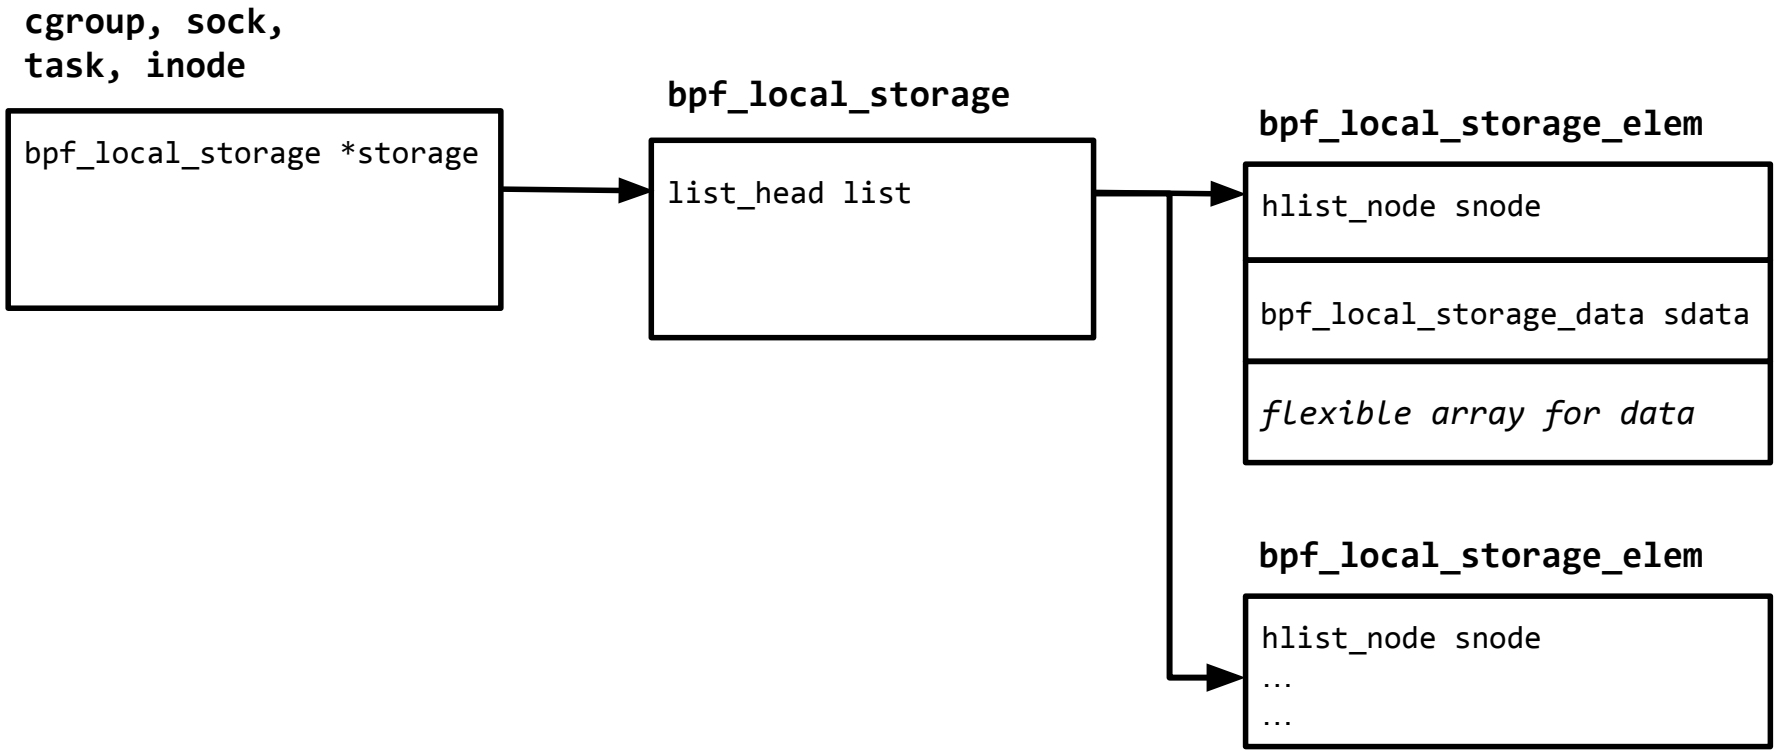

BPF local storage

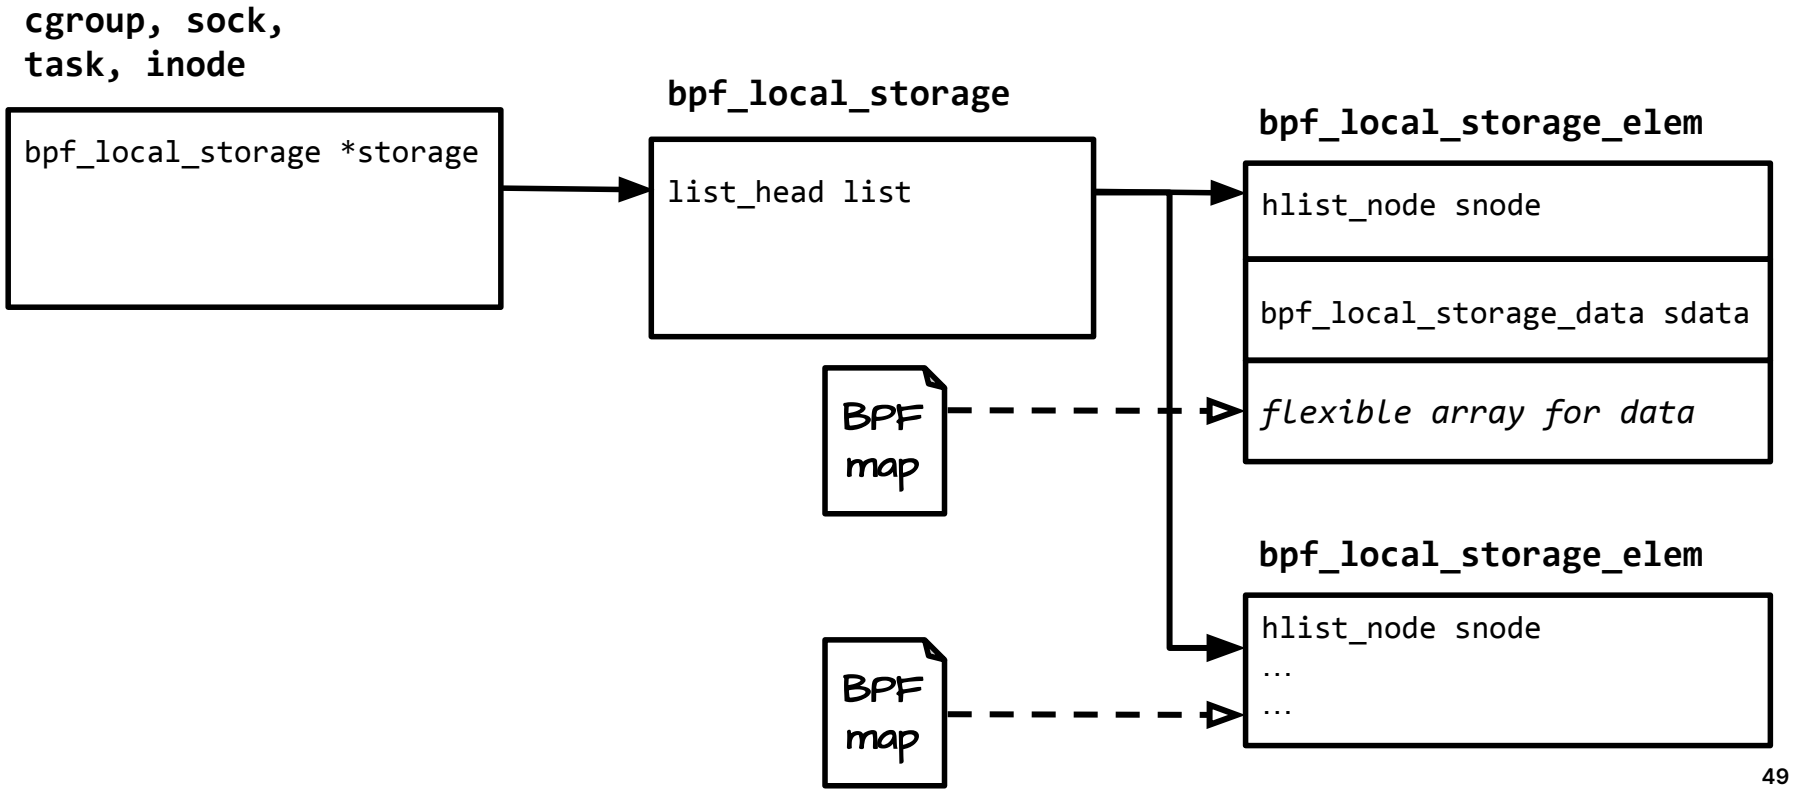

skb extension for BPF local storage

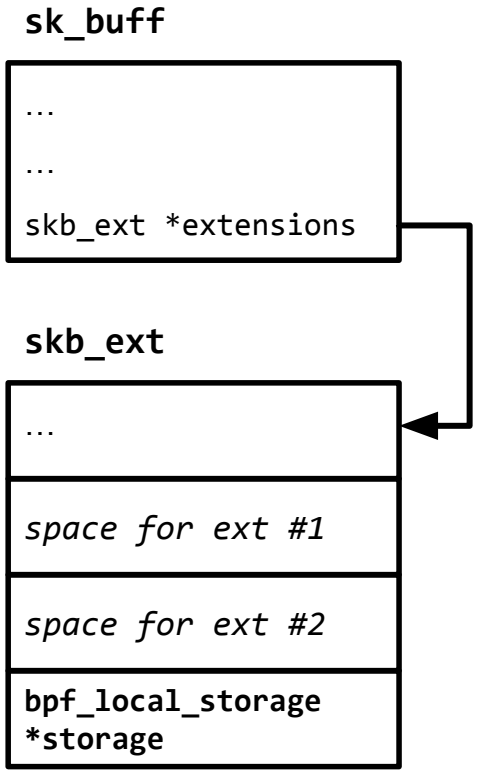

skb extension for BPF local storage

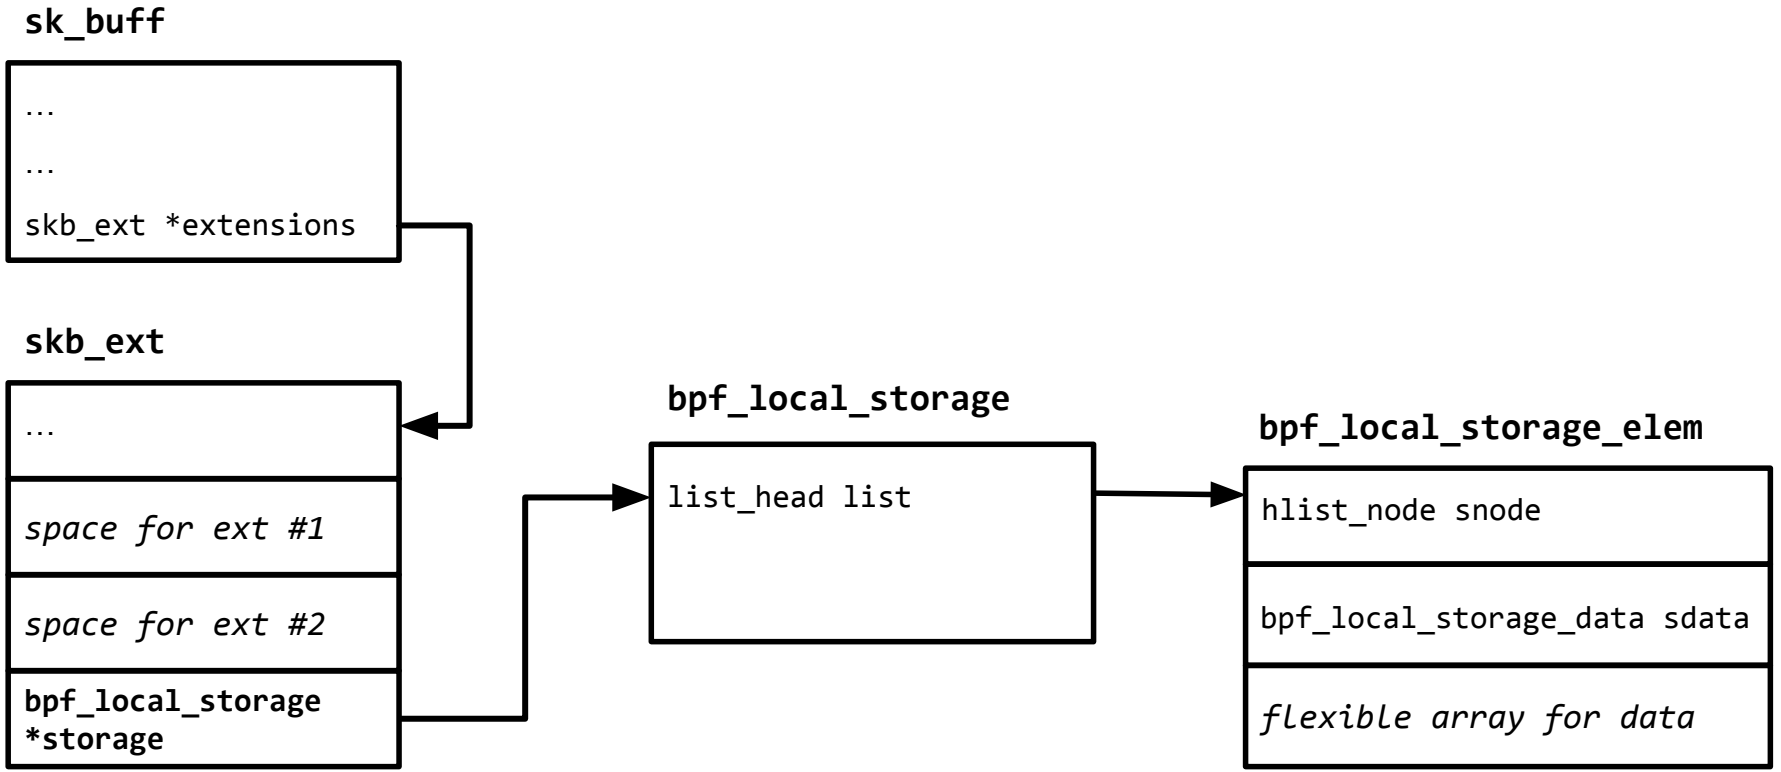

skb extension for BPF local storage

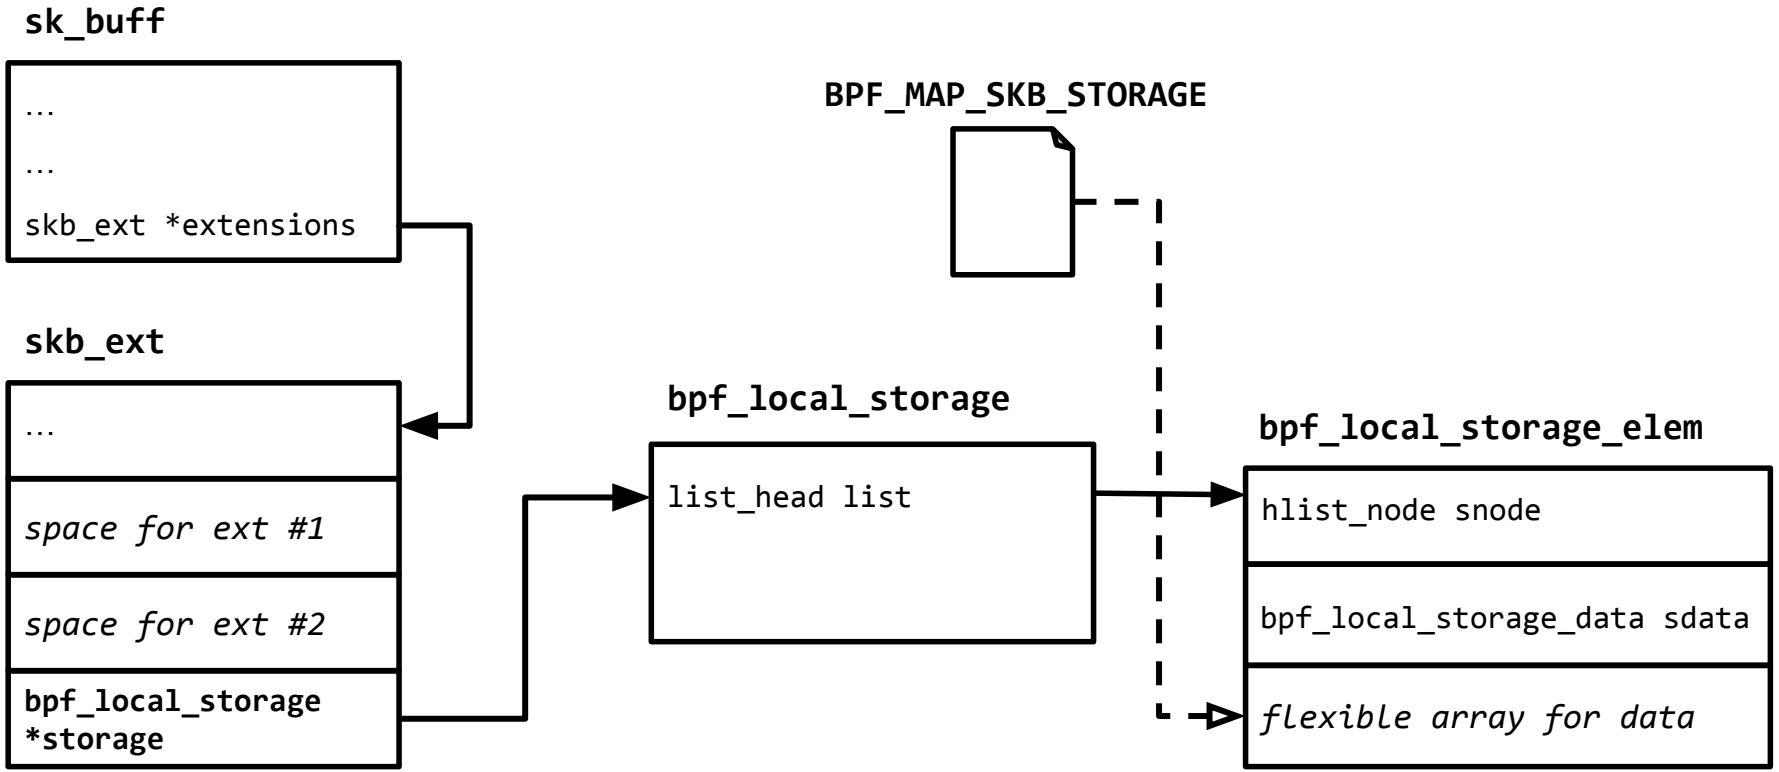

That’s all neat but...

skb extension for BPF local storage

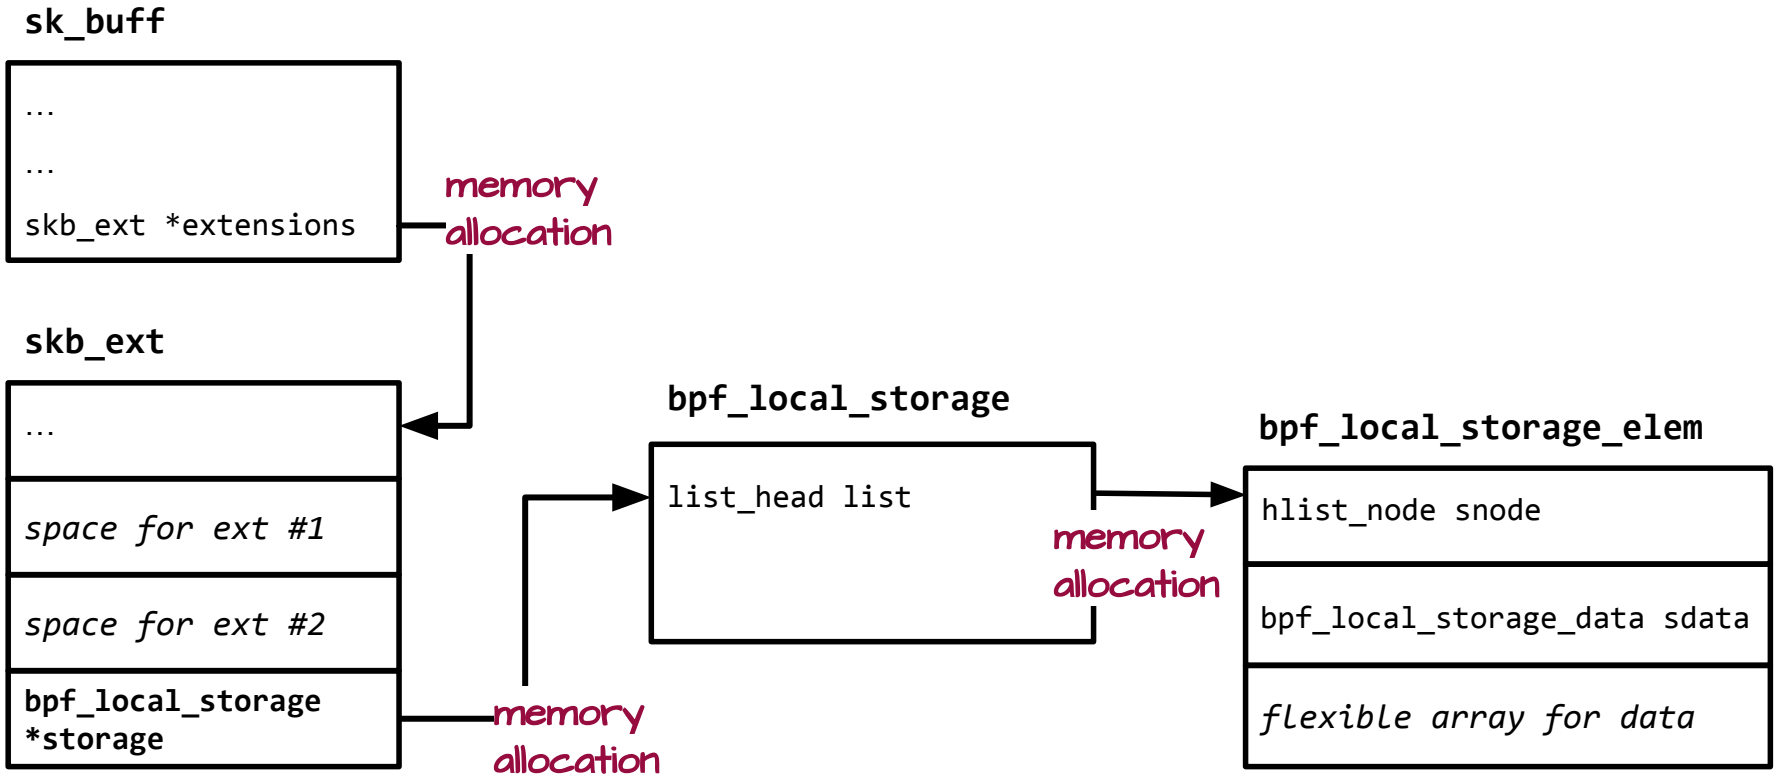

AI-assisted prototype for BPF local storage skb extension

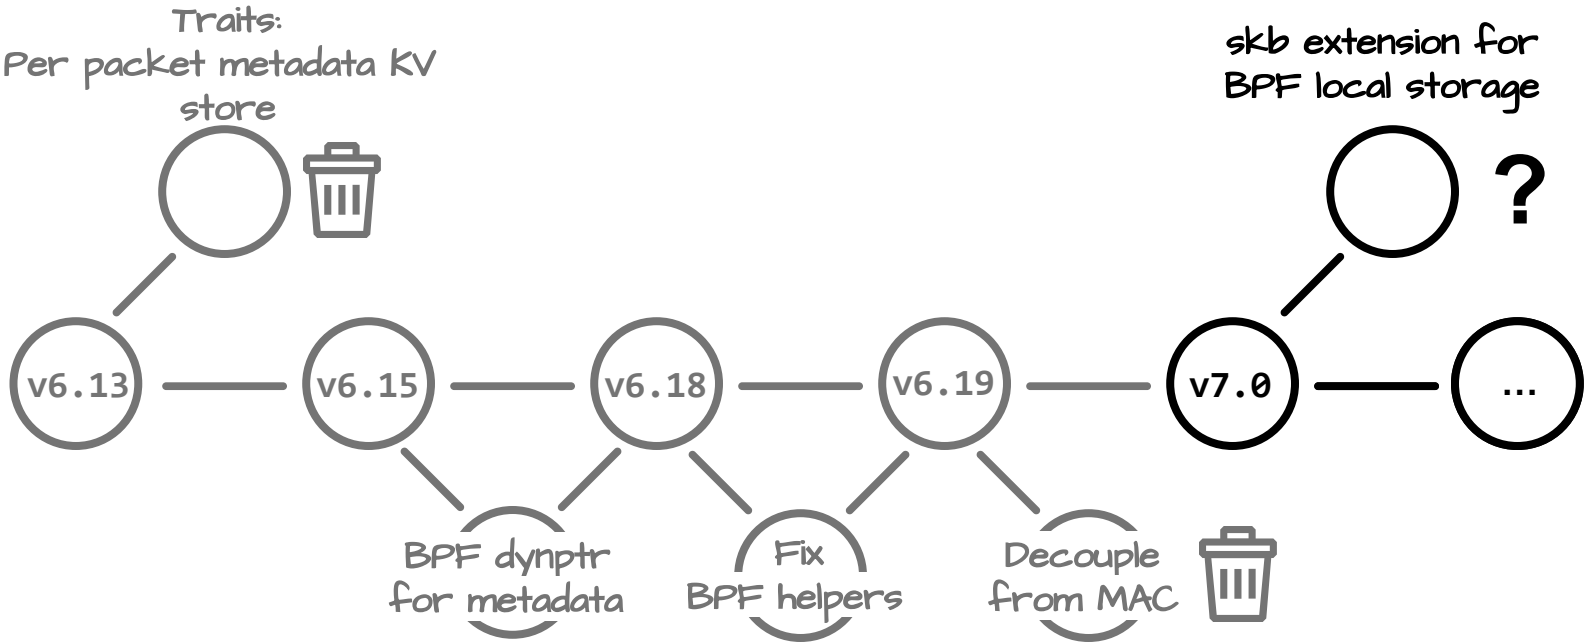

Make XDP / skb metadata work

CLOUDFLARE

Third time's the charm? Time will tell... 🤞

Resources (newest to oldest)

CLOUDFLARE

Dev branches

1. AI-assisted prototype for BPF local storage skb extension

Patch series

1. [PATCH RFC bpf-next 00/15] Decouple skb metadata tracking from MAC header offset
2. [PATCH bpf-next v4 00/16] Make TC BPF helpers preserve skb metadata
3. [PATCH bpf-next v7 0/9] Add a dynptr type for skb metadata for TC BPF
4. [PATCH RFC bpf-next v2 00/17] traits: Per packet metadata KV store

Talks

1. Linux Plumbers 2025: Packet Metadata - Where Are Thee?
2. Netdev 0x19 2025: Traits: Rich Packet Metadata
3. Linux Plumbers 2024: Marking Packets With Rich Metadata

Thank you

jakub@cloudflare.com

A bunch of question marks, Wikimedia Commons, CC BY-SA 4.0 DEED

Cloudflare’s goal to hire 1,111 interns in 2026

https://blog.cloudflare.com/cloudflare-1111-intern-program/