KRISTOF PROVOST

A PACKET'S JOURNEY THROUGH PF

WHO AM I?

- Kristof Provost
- kp@FreeBSD.org
- pf (in FreeBSD) maintainer since 2015
  - “Hmm, IPv6 fragmentation handling isn’t great. I bet I could fix that!”
  - And in pfSense since 2021
    - Thanks, Netgate!

INTRODUCTION

- Based on FreeBSD main as of today(-ish)
- See also “A Packet’s Journey Through the OpenBSD Network Stack”
  - Alexander Bluhm
  - https://www.youtube.com/watch?v=Kn2XEW4Qre0
  - https://2024.eurobsdcon.org/slides/eurobsdcon2024-alexander_bluhm-a_packets_journey.pdf

TL;DR: THE NETWORK STACK

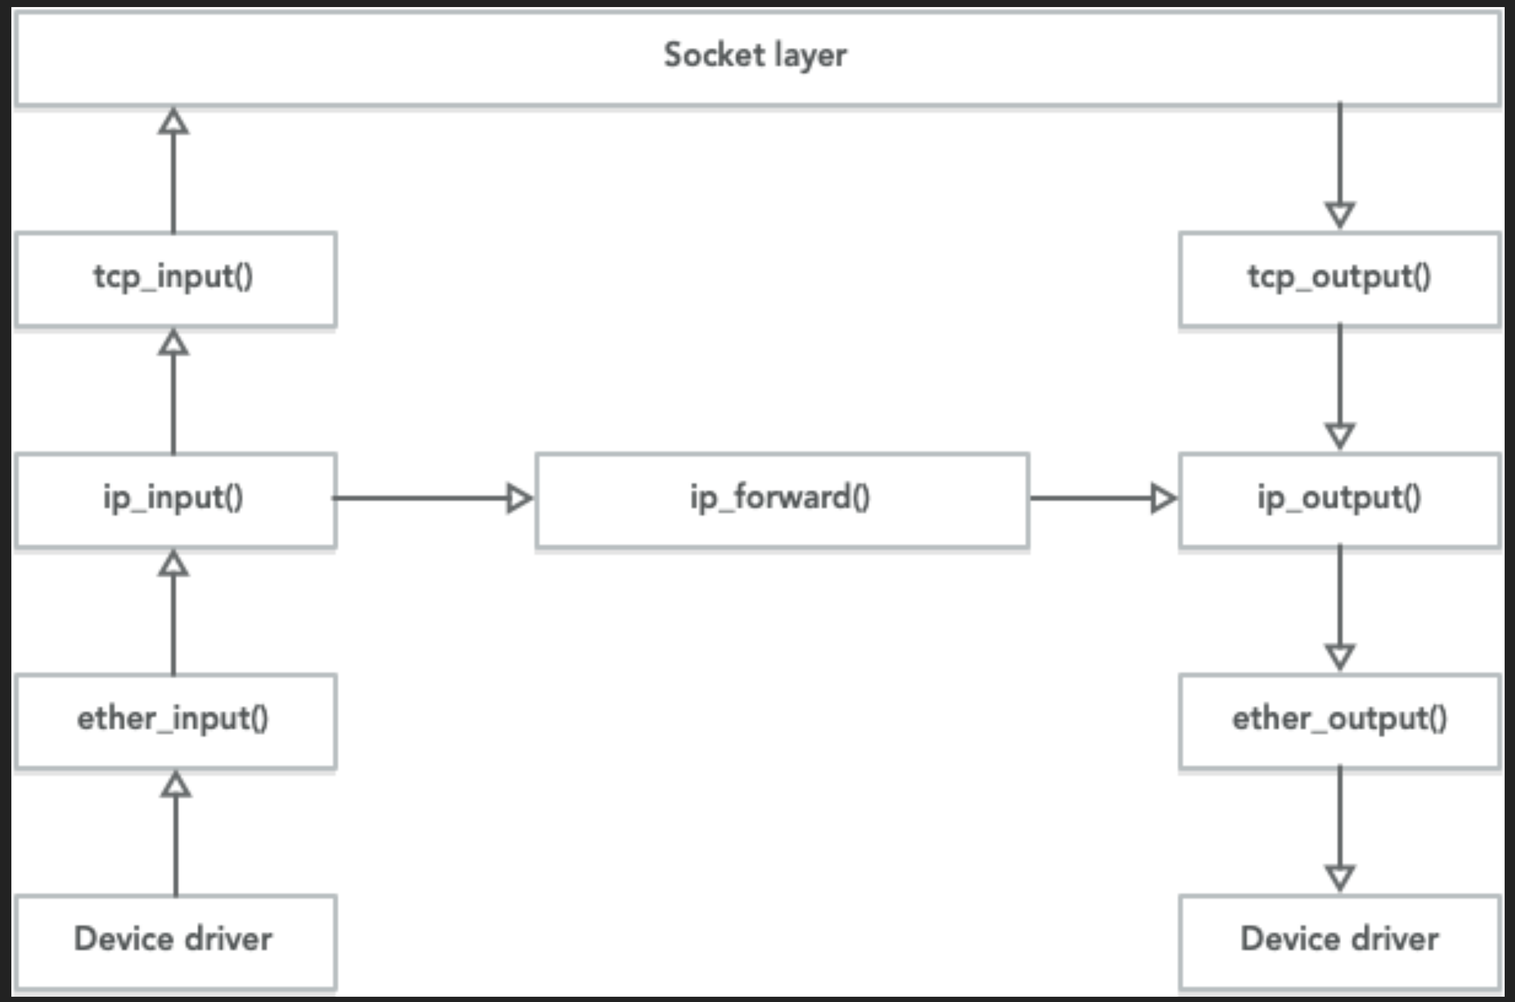

TL;DR: THE NETWORK STACK

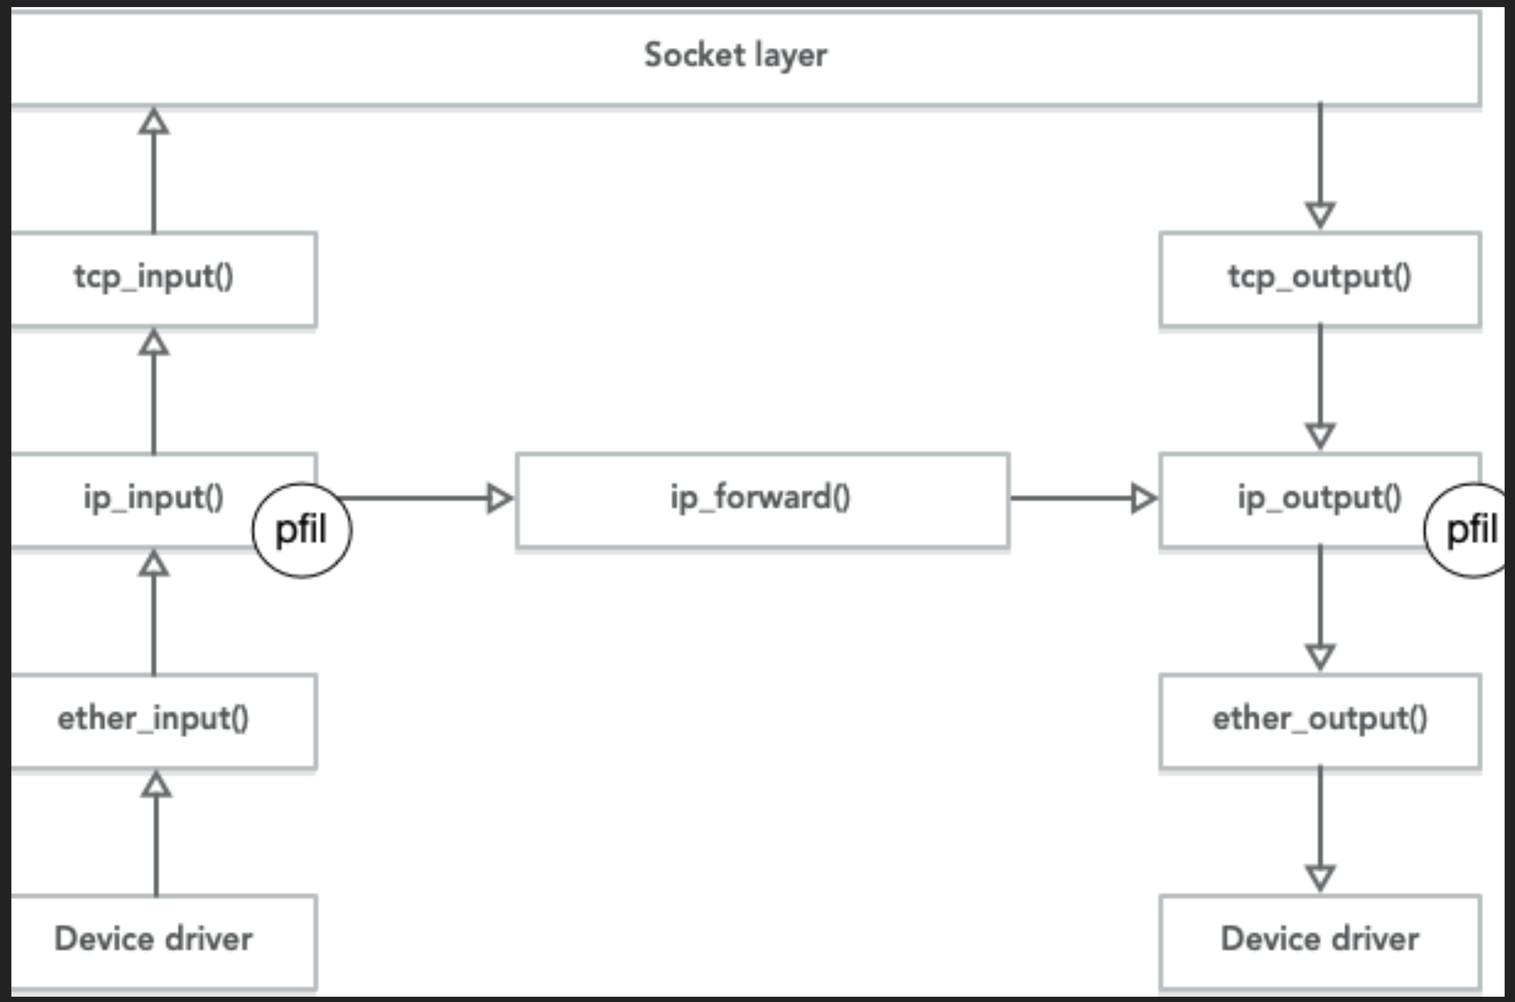

KEY CONCEPTS

- States
  - pf is a stateful firewall
  - Even for stateless protocols (I.e. UDP)
- Rules
  - i.e. what policy are we apply to packets (or connections!)

30,000 FT OVERVIEW

- pf_test()
- pf_setup_pdesc()
  - Parse packet
  - Normalise packet
    - i.e. reassembly
- pf_test_state_<protocol>()
  - (TCP, UDP, SCTP, ICMP, Other)
  - Find state
  - Or pf_test_rule()

30,001 FT OVERVIEW

- Output handling
  - pass
  - drop
  - route-to
  - af-to
- IPv6 special case
  - Re-fragment

IMPLICATIONS

- Test for state first
- Evaluate rules only if no state is found

- So if rules change, existing connections keep passing
  - ‘block all’ may not be block everything immediately!
  - Flush or kill states to actually terminate them

MORE IMPLICATIONS

- State lookup is performance critical
- How does this work?
  - Hash table
    - With linked list of states in each hash row
      - net.pf.states_hashsize
  - Key
    - Src/dst IP
    - Src/dst port (or ICMP type/code)
    - Address Family
    - Protocol

CONTROL PLANE

- How the user configures pf and get information out of it
- Interface to userspace
  - ioctl
  - ioctl + nvlist
  - netlink
    - Hopefully the only option in the future
- Somewhat abstracted by libpfctl
- pfctl

QUESTIONS?

SPARE SLIDES

ETHERNET

- FreeBSD-unique feature
- (Very) basic filtering on Layer 2
  - Mostly so we can look at MAC addresses for captive portal scenarios
- Stateless

SCTP

- Very TCP-like, but with multiplexed flows
- And multihoming
- Hence special case handling
- Parse SCTP header to find ASCONF chunks
- Set up states for all multi homed options

DUMMYNET

- Traffic shaping
-