Faster OpenBSD Networking

Alexander Bluhm

bluhm@openbsd.org

EuroBSDCon, September 2022

Agenda

Network Layers

- syscall
- file descriptor
- socket
- IP protocol
- IPsec
- IP input, output, forwarding
- pf
- routing, ARP, ND6
- interface
- driver
- hardware

Other Layers

- interrupts
- malloc, pools
- tasks
- multicast
- ifconfig ioctl
- pseudo devices: aggr bpe bridge carp egre enc eoip etherip gif gre lo mgre mpe mpip mpw nvgre pair pflog pflow pfsync ppp pppoe svlan tap tpmr trunk tun veb vether vlan vport vxlan wg

Developement

- ideal
  - make it MP safe
  - run in parallel
  - make it fast
  - advance in steps
- reality
  - work on whatever makes fun
  - deal with 40 years old code

History

- interrupts with spl
- kernel lock
- kernel threads
- MP safe subsystems
- netlock
- multiqueue NIC
- shared netlock

Single Processor

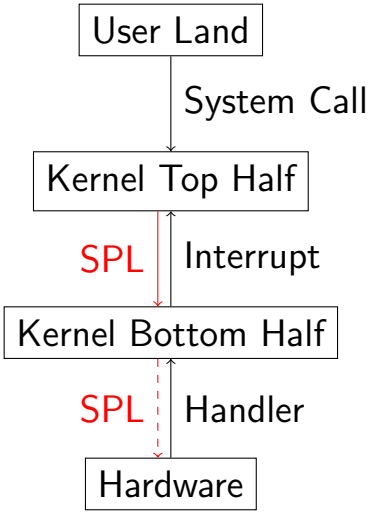

Multi Processor Kernel Lock

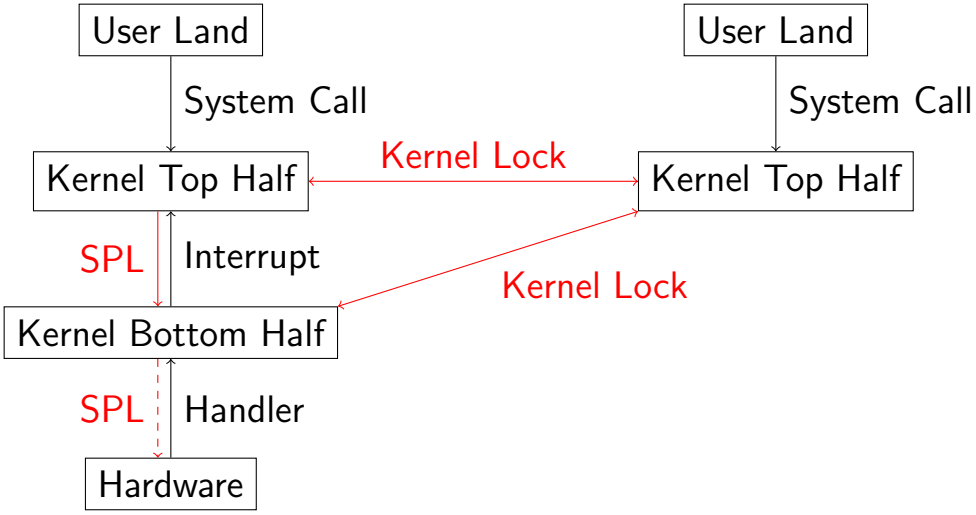

Multi Processor Fine Grained Mutex

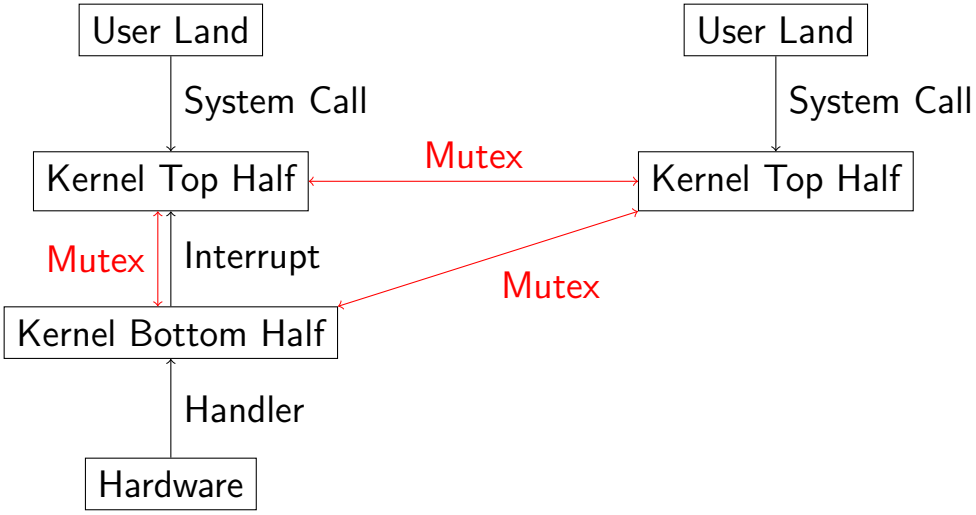

Multi Processor Multi Queue

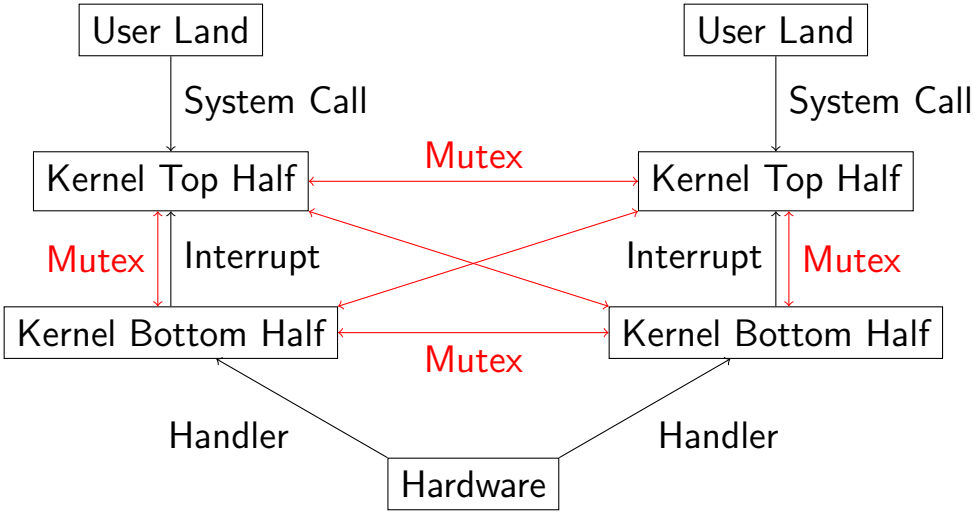

Kernel Threads

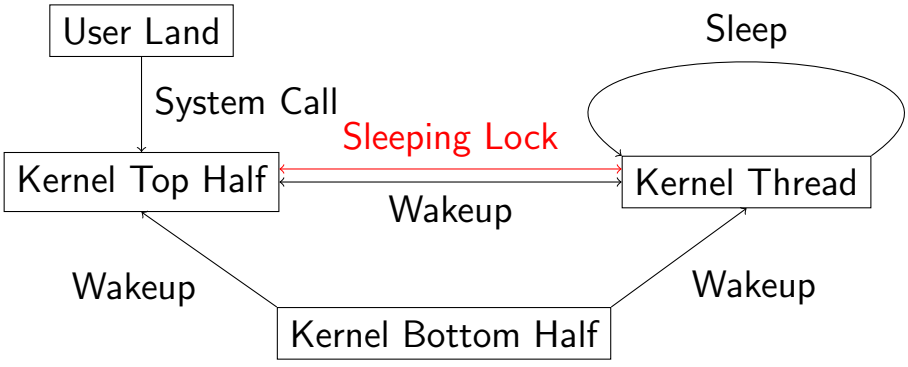

Softnet Receive

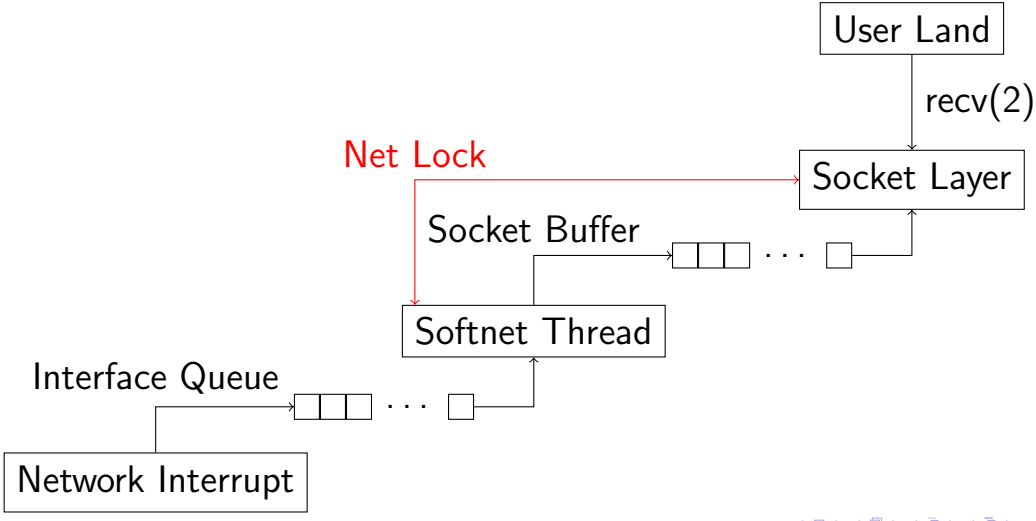

Softnet Send

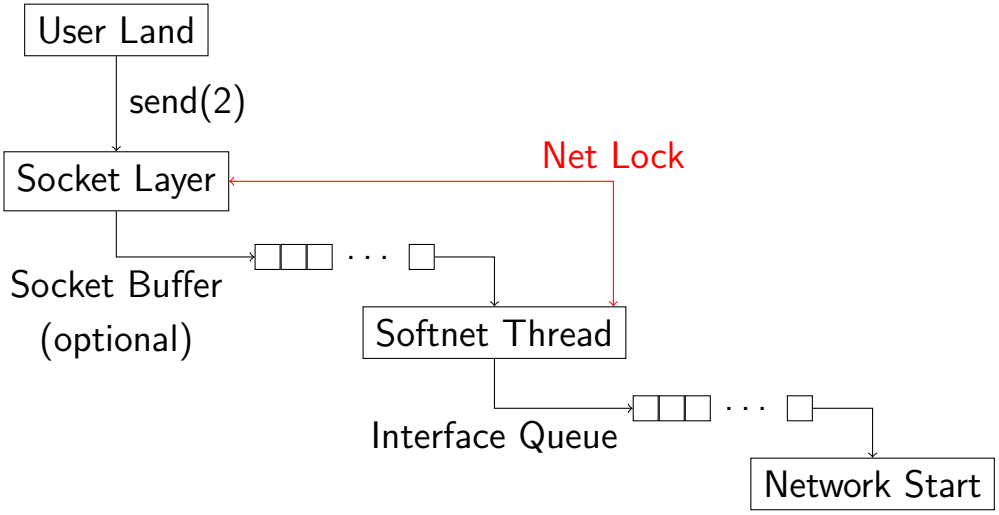

Softnet Forwarding

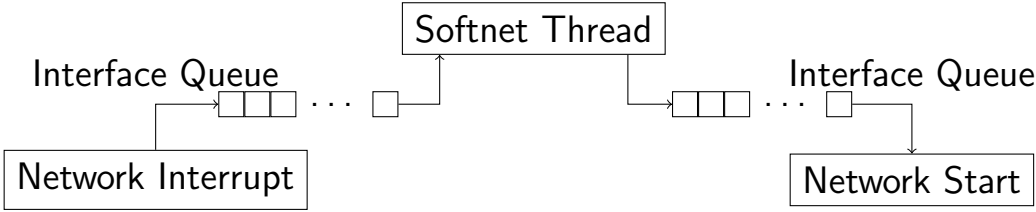

Configuration Lock

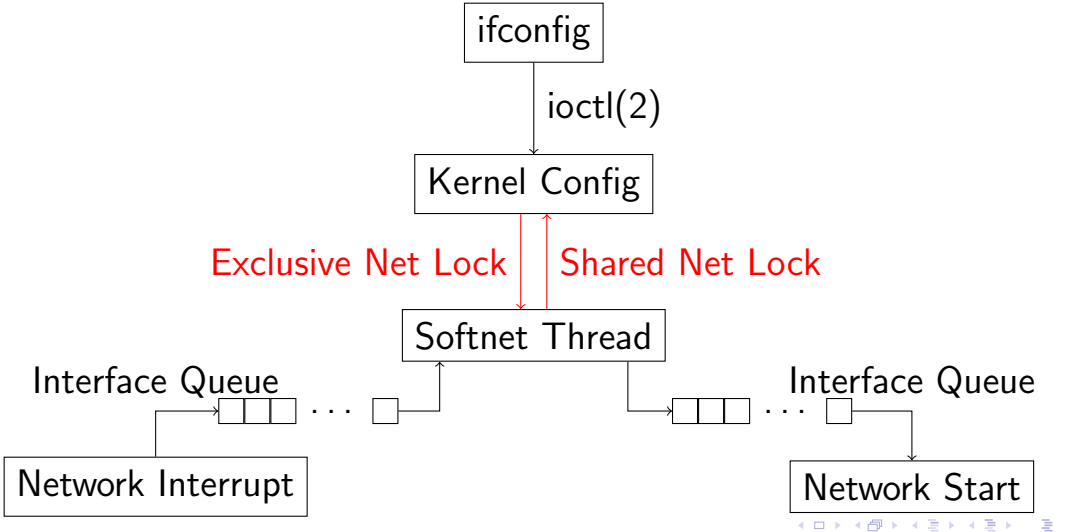

Parallel Forwarding

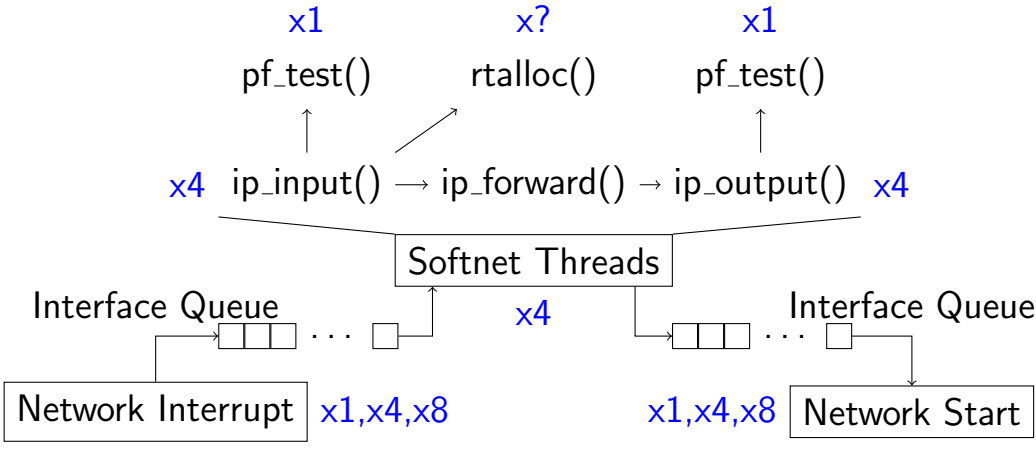

Protocol Queue

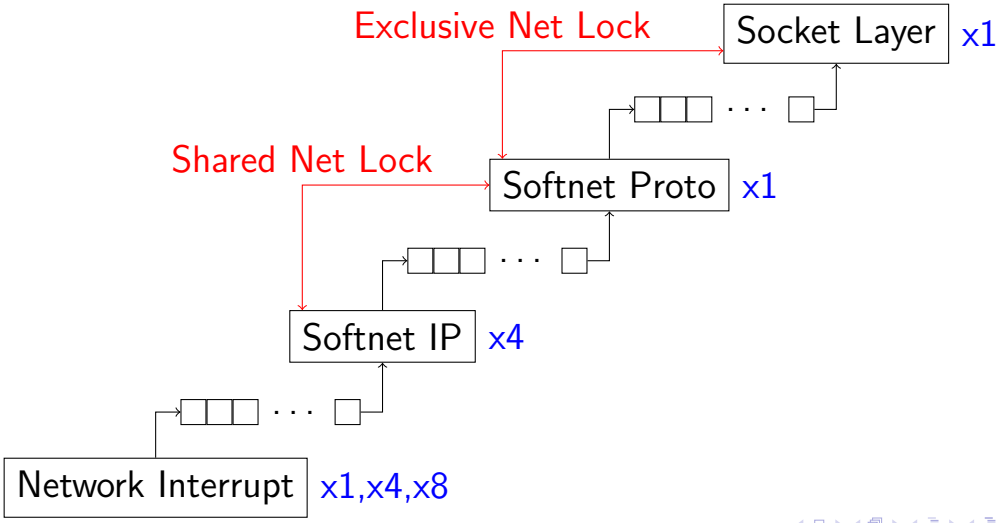

Parallel Receive, UDP coming soon

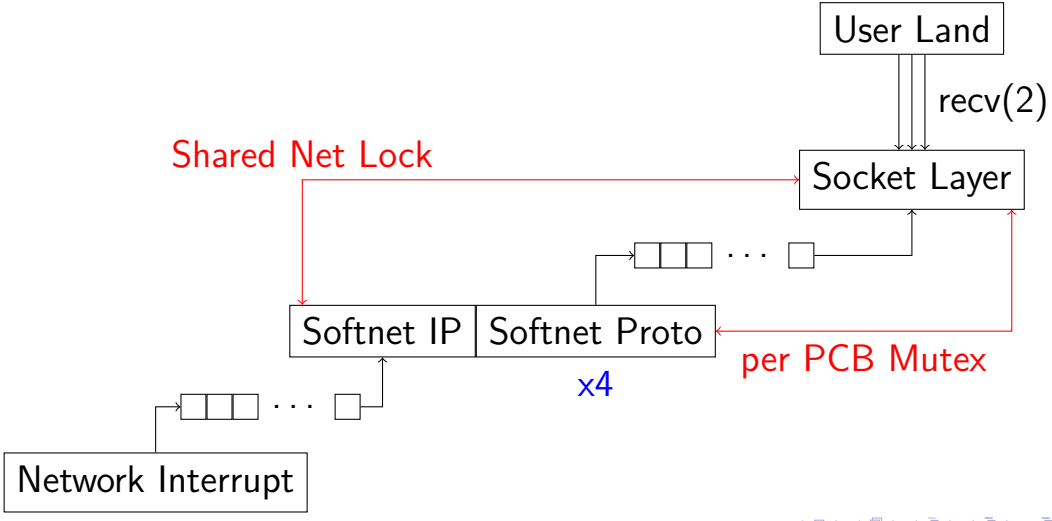

Hardware Setup

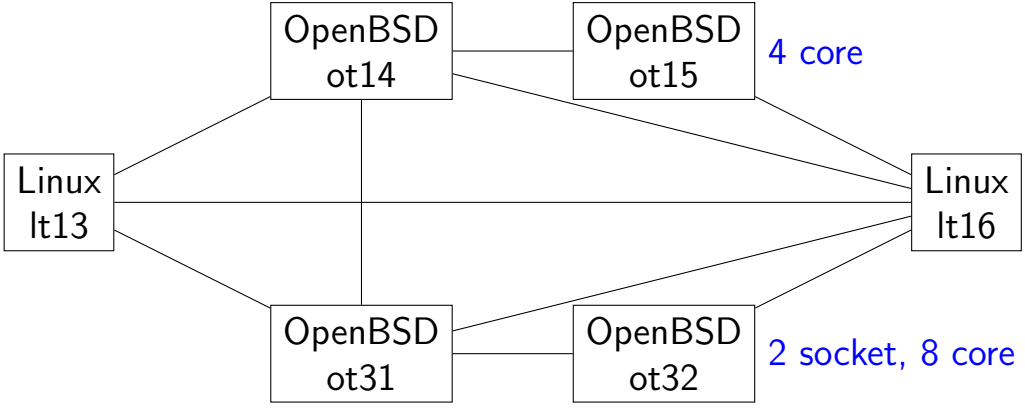

7.1 Forwarding

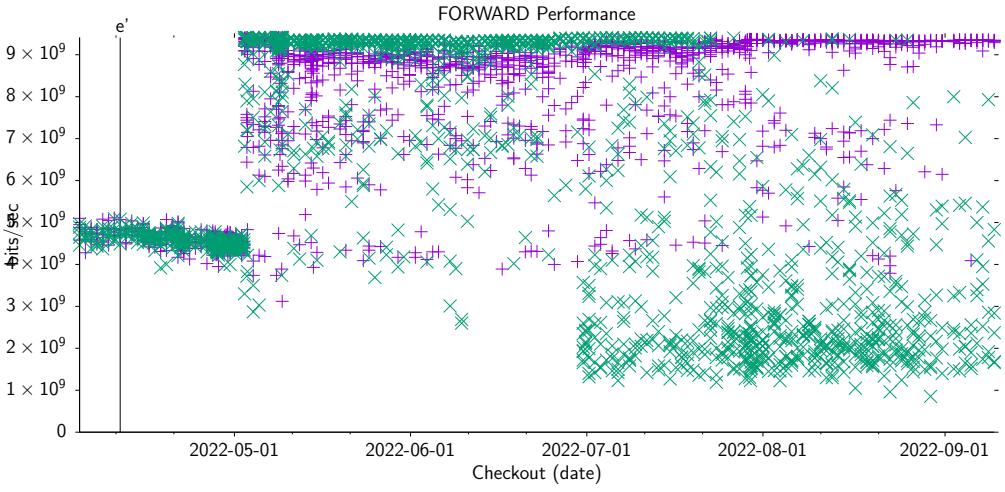

7.1 TCP Send Receive

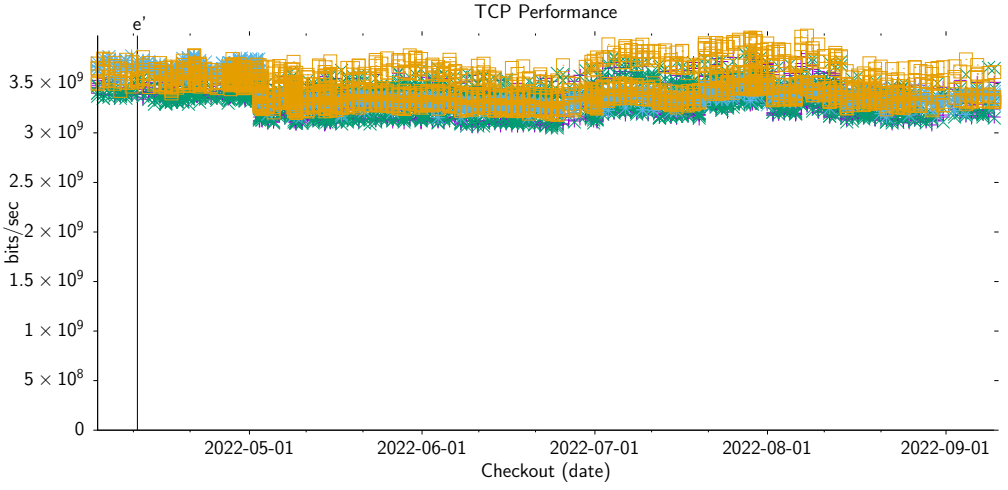

7.1 TCP IPv6 Send Receive

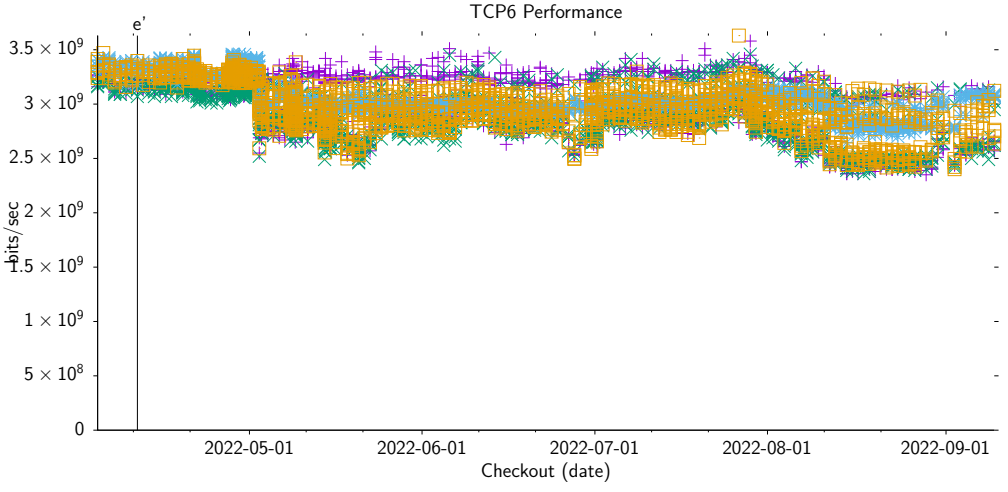

6.9 IPsec

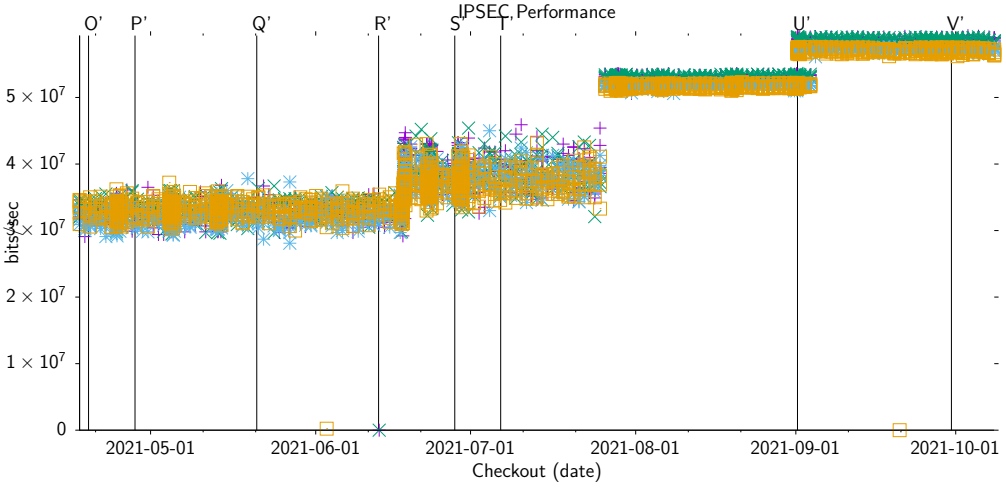
IPSEC, Performance
O'
P'
Q'
R'
S'
T'
U'
V'
5 × 10^7
4 × 10^7
3 × 10^7
2 × 10^7
1 × 10^7
0
bits/sec
2021-05-01
2021-06-01
2021-07-01
2021-08-01
2021-09-01
2021-10-01
Checkout (date)

UDP Receive

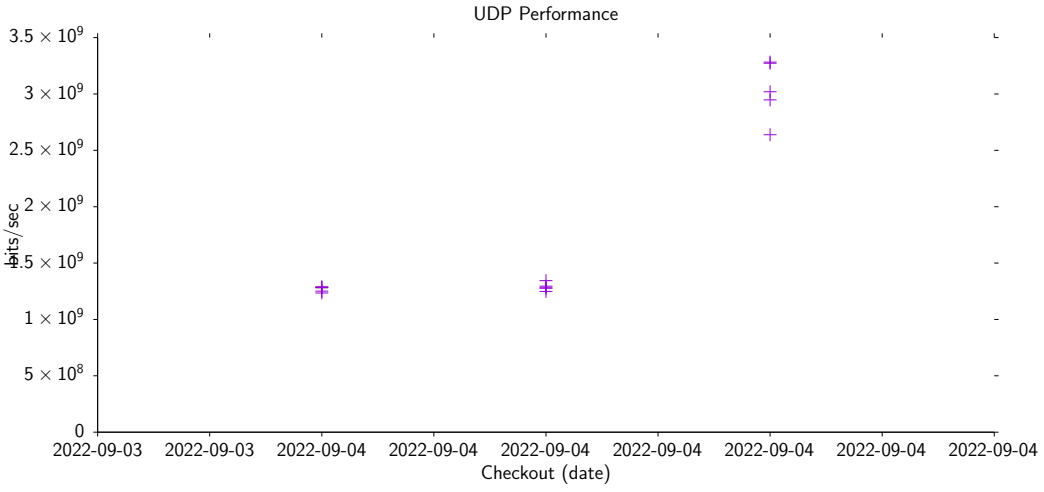
UDP Performance
3.5 × 10^9
3 × 10^9
2.5 × 10^9
2 × 10^9
1.5 × 10^9
1 × 10^9
5 × 10^8
0
bits/sec
2022-09-03
2022-09-03
2022-09-04
2022-09-04
2022-09-04
2022-09-04
2022-09-04
2022-09-04
2022-09-04
Checkout (date)

Results UDP Parallel

http://bluhm.genua.de/perform/results/2022-09-04T20:23:45Z/perform.html

Flamegraph UDP Receive

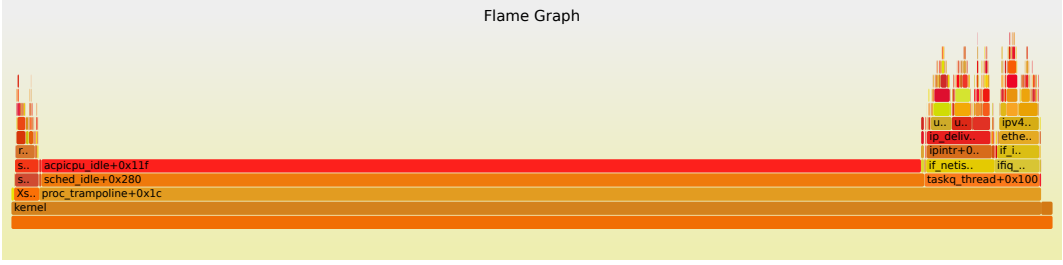

Flamegraph UDP Receive Parallel

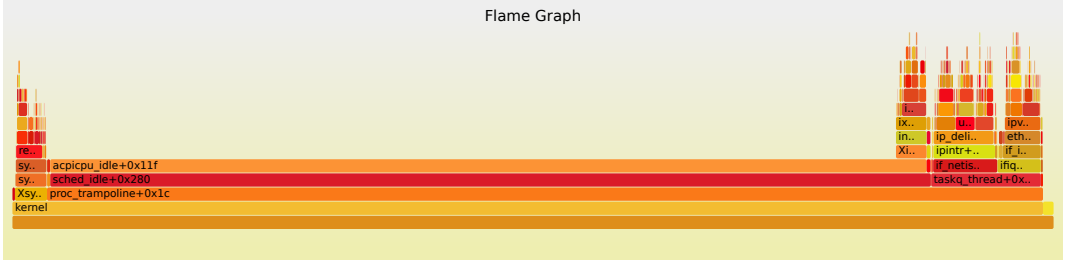

Flamegraph UDP Parallel

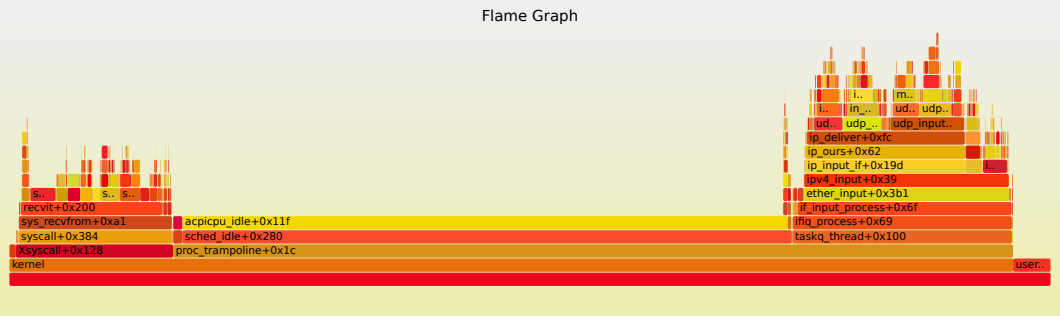

UDP Recv 8 Netlock

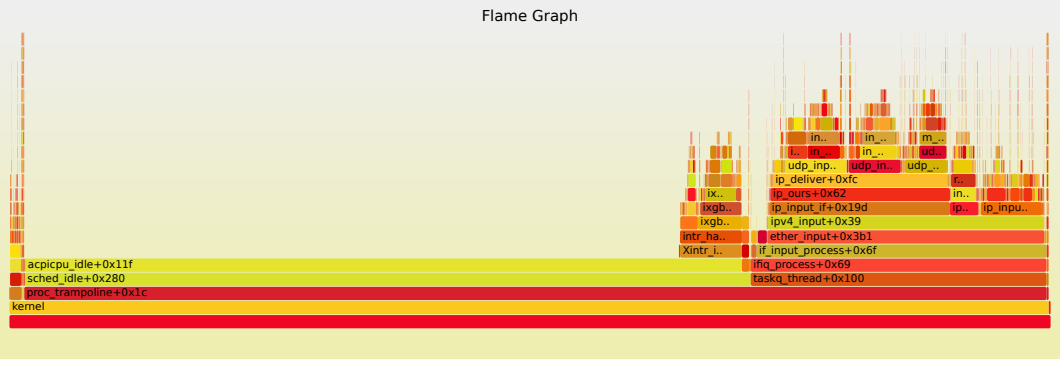

UDP Send 8 Netlock

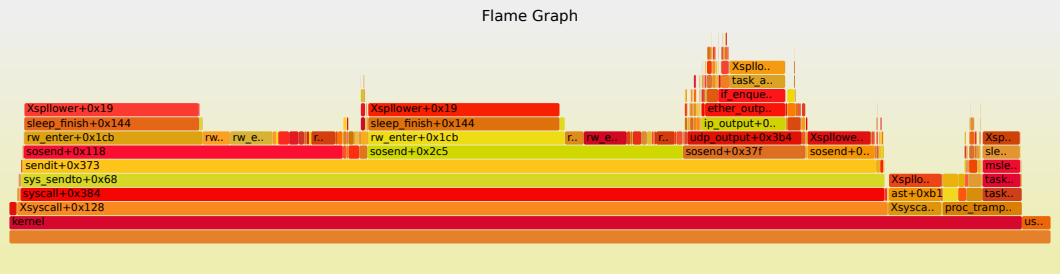

Links

- http://bluhm.genua.de/
- http://bluhm.genua.de/perform/results/perform.html
- http://bluhm.genua.de/perform/results/test.data
- https://github.com/bluhm/regress-all
- https://github.com/bluhm/udpbench
- https://github.com/bluhm/talk-networkmp

Questions