genua.

Update on OpenBSD Networking Performance Improvements

Alexander Bluhm

EuroBSDCon, September 2025

genua.

Agenda

01 History

02 Packet Processing

03 Performance

04 Speed

Single Processor SPL, 1980s

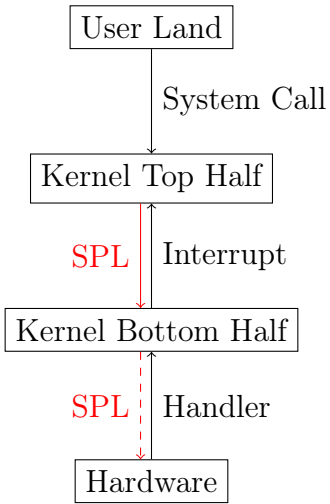

Softnet AST, 1990s

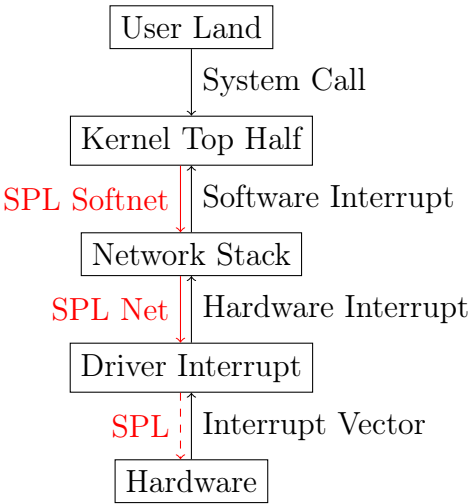

genua.

Multi Processor Kernel Lock, 2000s

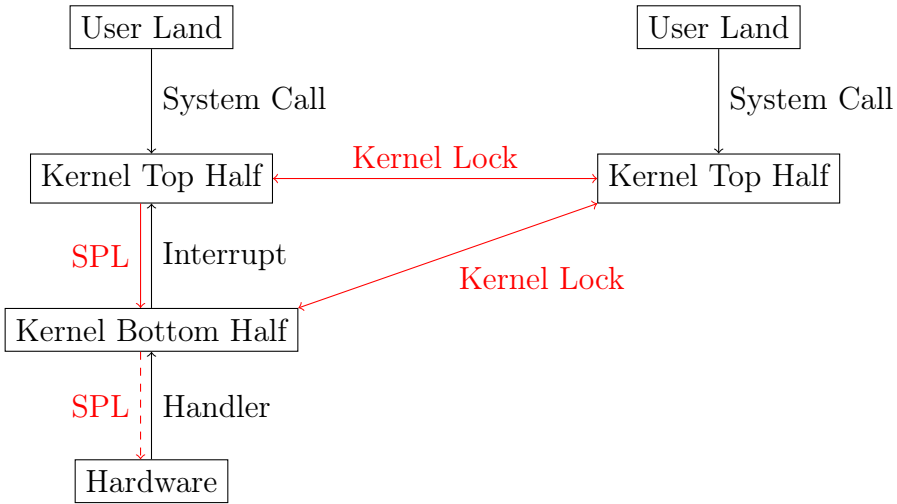

genua.

Kernel Lock and Net Lock, 2015

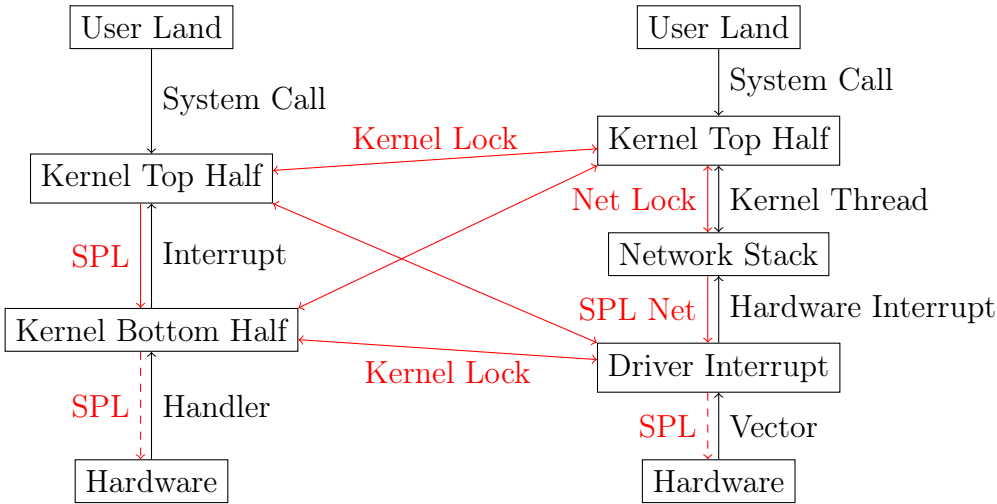

genua.

Unlock Network System Calls, 2018

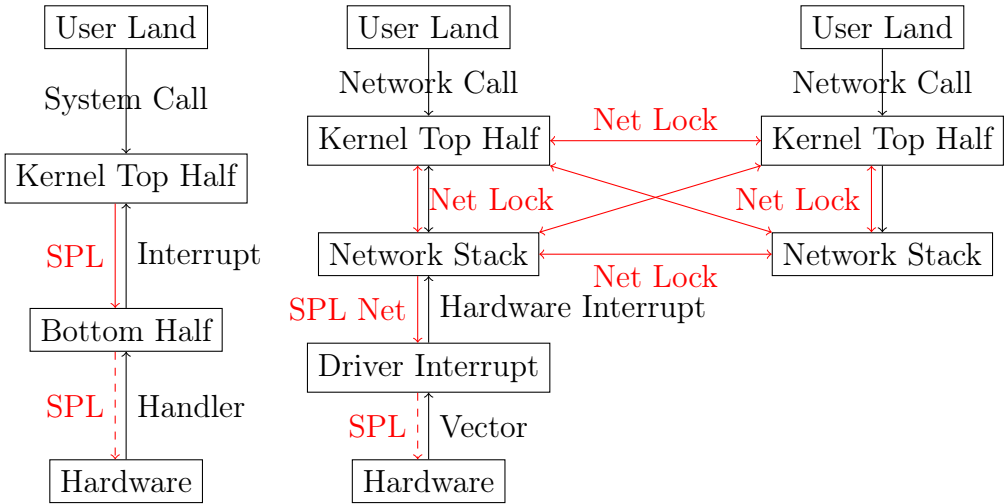

genua.

Multi Queue Drivers, 2020

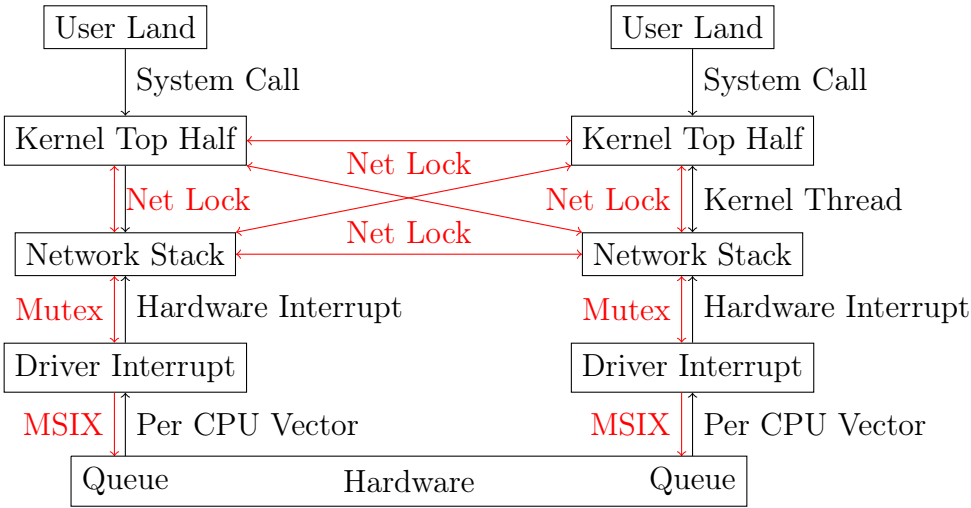

genua.

Unlock Socket Receive and Send, 2022 and 2024

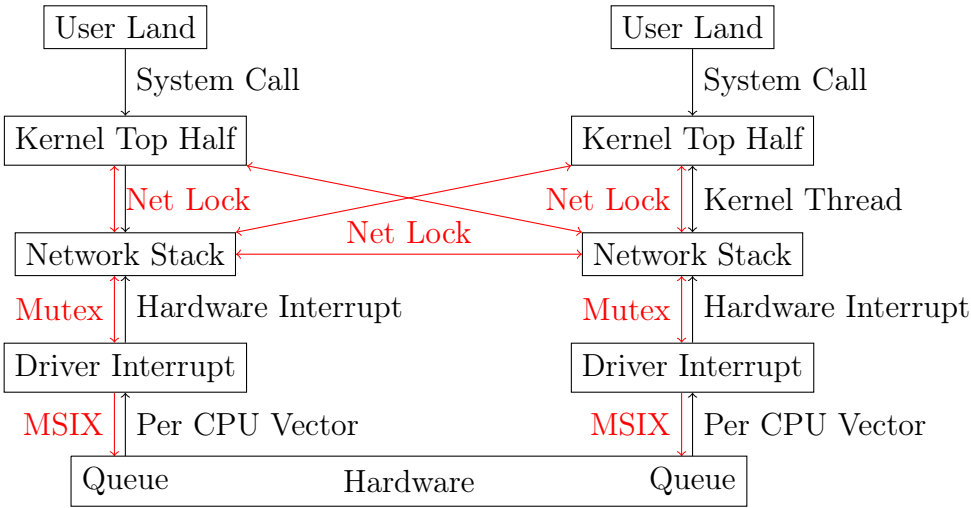

genua.

Parallel Forwarding, 2022

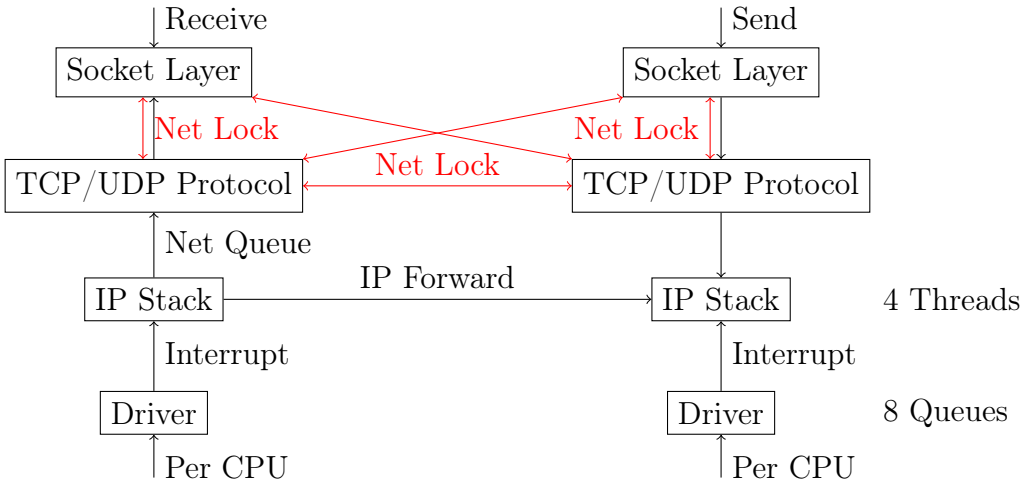

Parallel UDP and TCP, 2024 and 2025

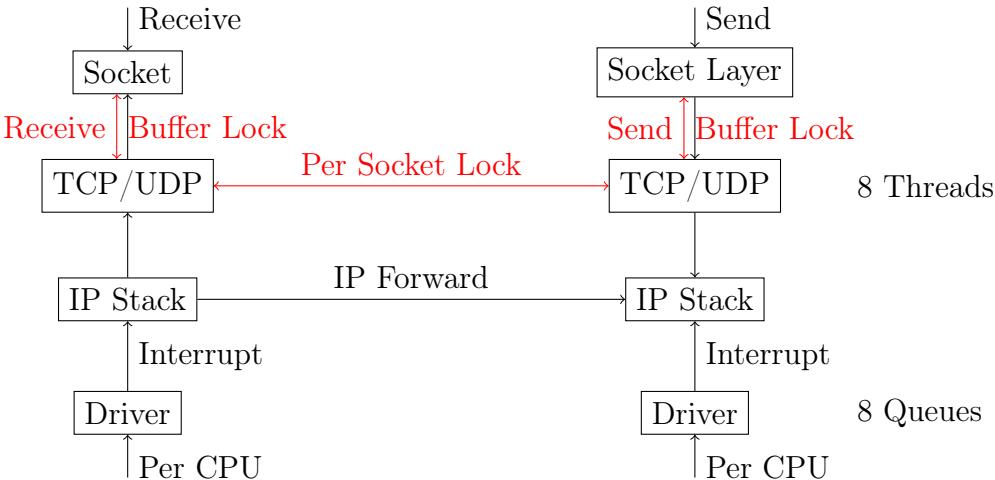

genua.

Good Mature Hardware, from 2011

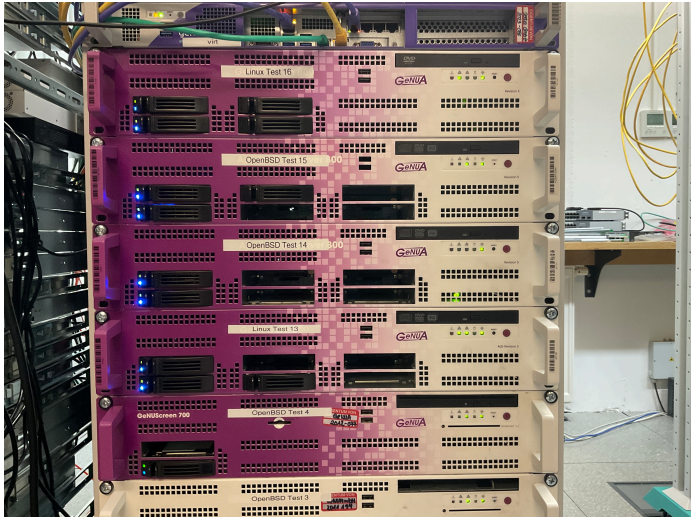

TCP Performance Testing for 7 Years

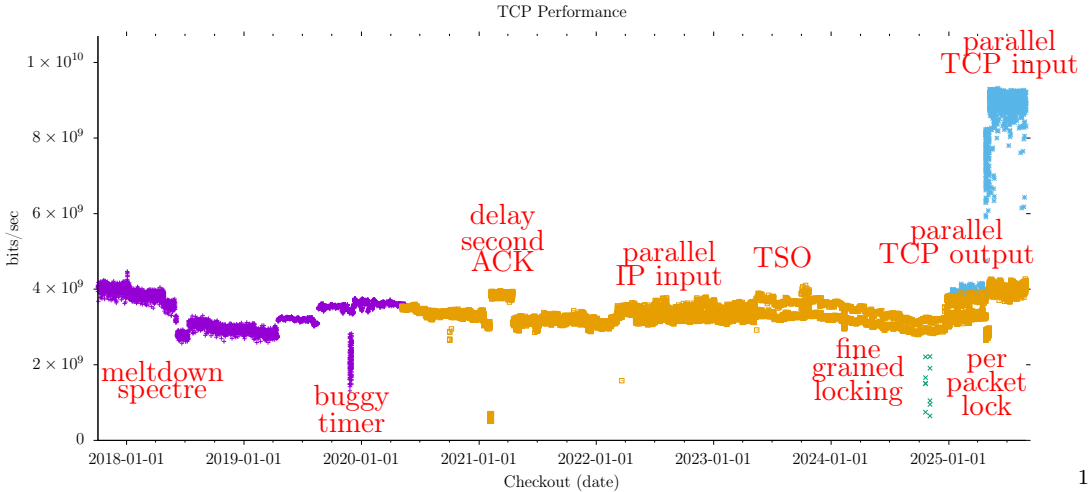

genua.

Network Protocol Stack

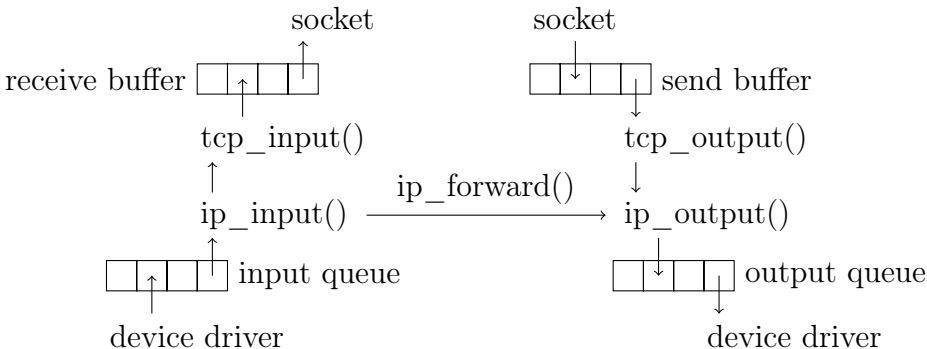

genua.

Towards Parallel Processing, 2024

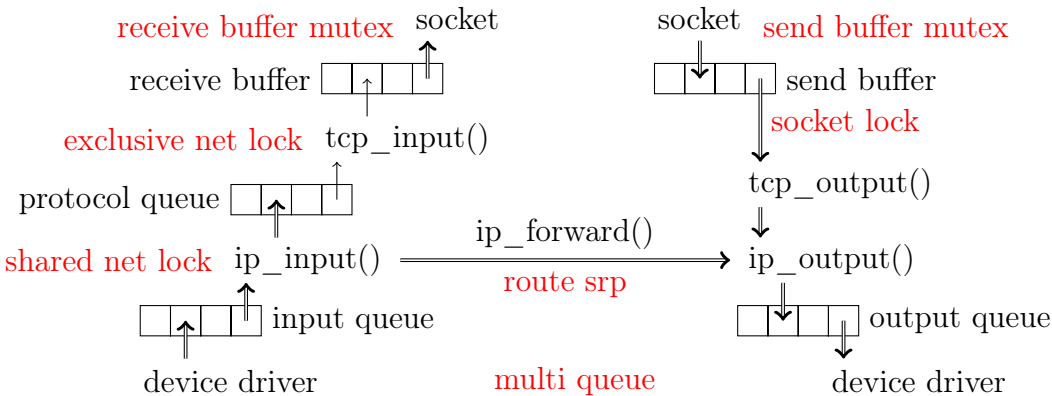

Parallel UDP Input, 2024

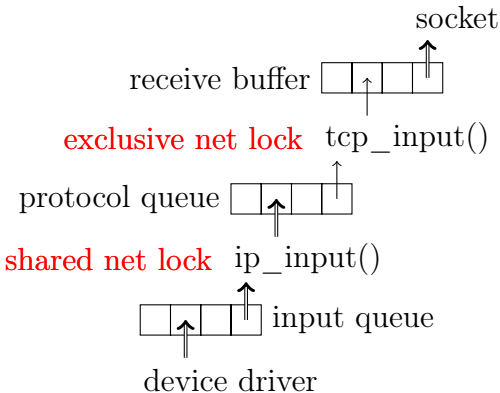

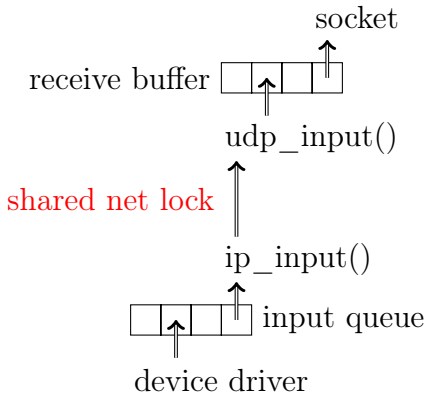

genua.

Parallel TCP Input, 2025

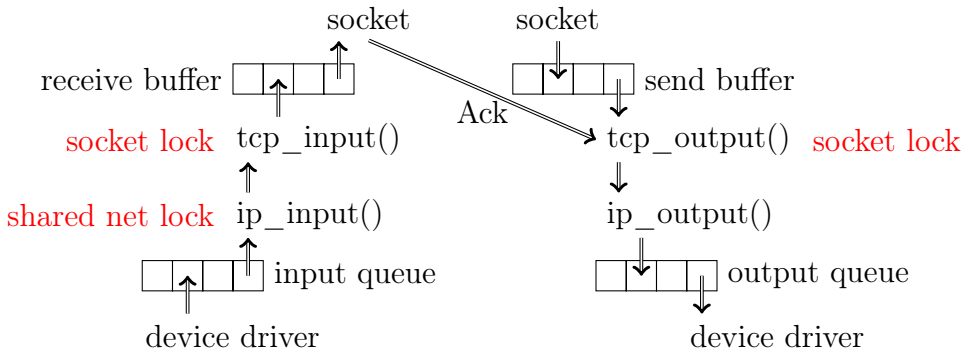

Exclusive TCP Receive Single Stream

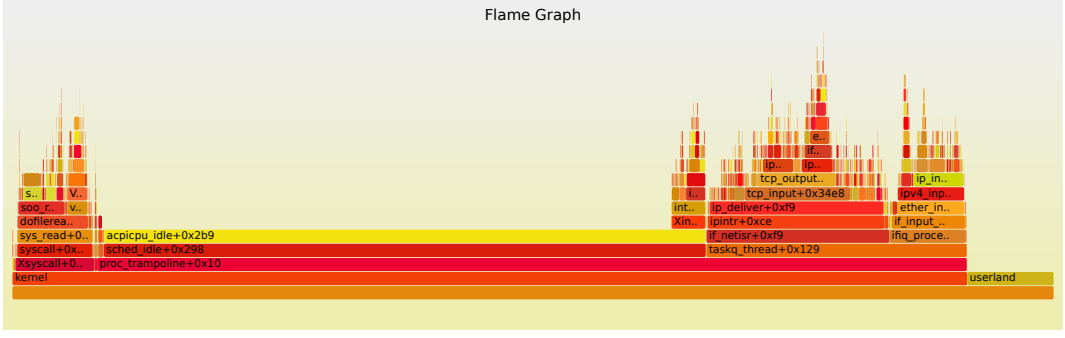

Parallel TCP Receive Single Stream

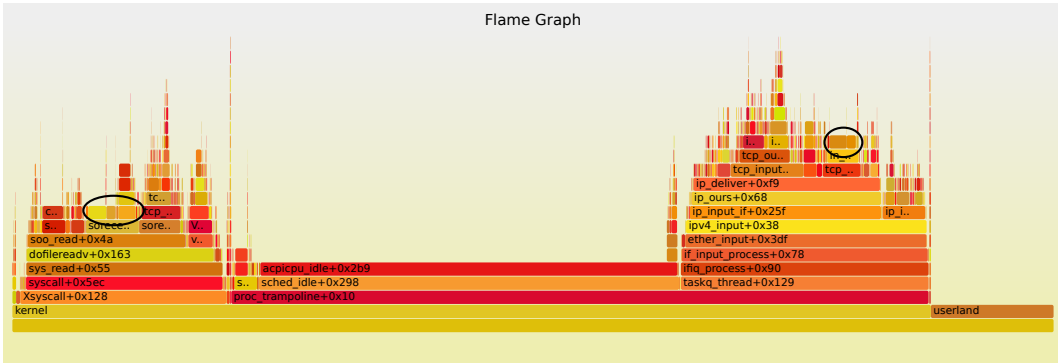

Parallel TCP Input, Socket Lock per Thread

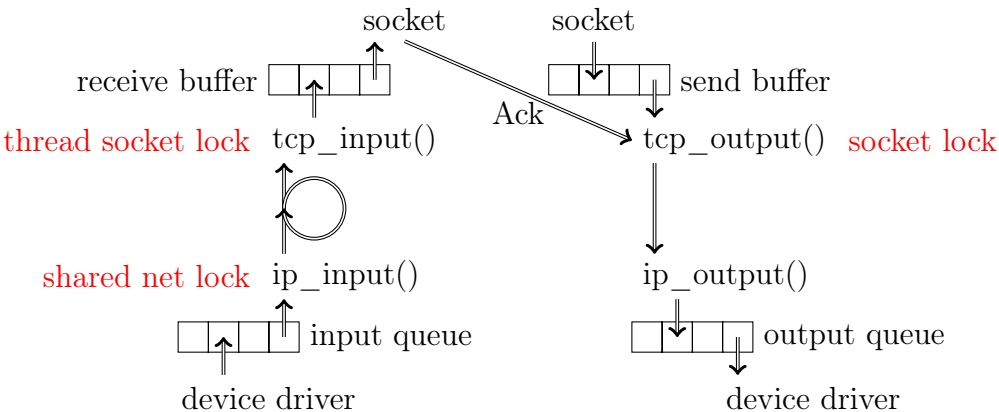

Variants for TCP Input

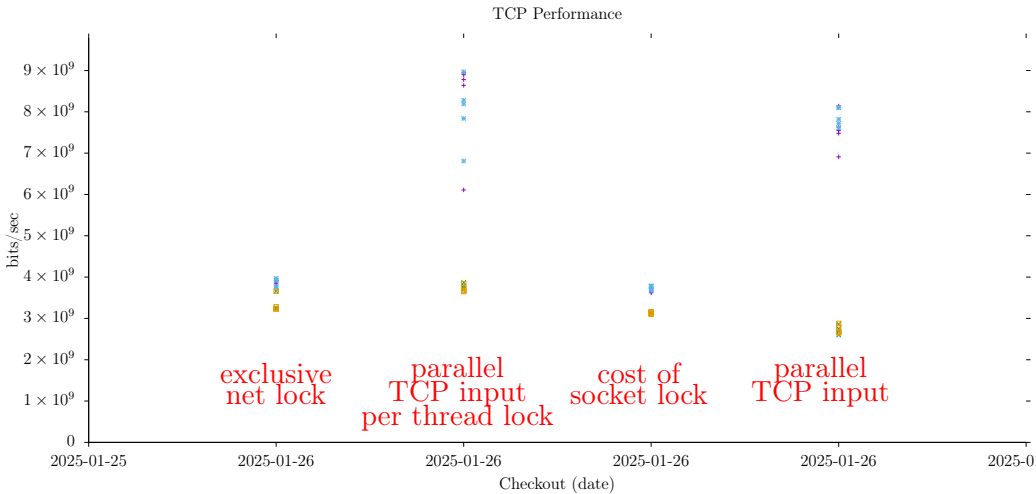

TCP Receive Parallel Stream, Socket Lock per Thread

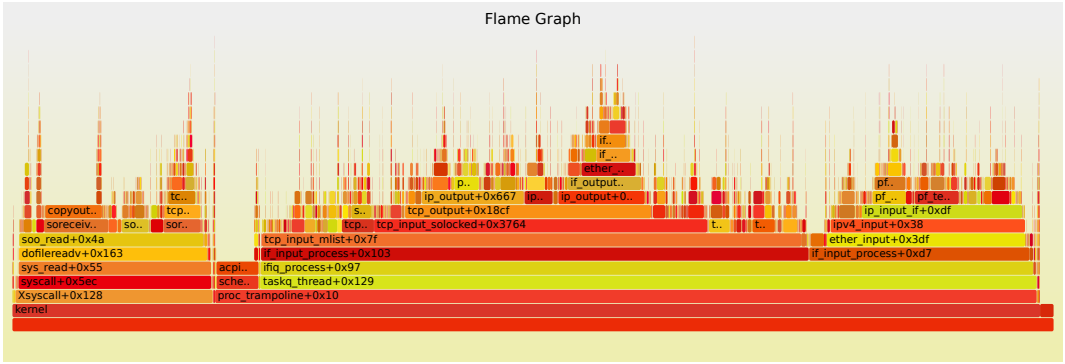

New 100 Gbit/sec Hardware, from 2024

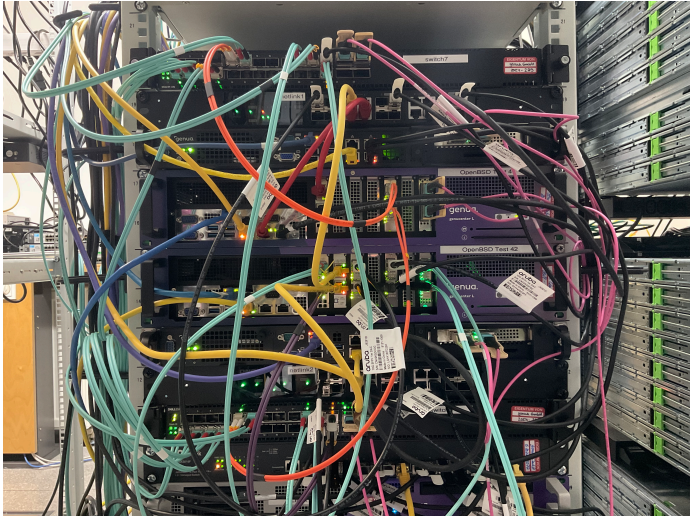

Daily Results

Top All Current Latest Running

OpenBSD netlink latest test results

created at 2025-09-03T13:39:35Z											
pass rate	**100%**	**100%**	**99%**	**100%**	**98%**	**99%**	**90%**	**98%**	**98%**	**90%**	
run at date	2025-09-03	2025-09-03	2025-09-03	2025-09-03	2025-09-03	2025-09-03	2025-09-03	2025-09-03	2025-09-02	2025-09-02	
sub runs	**1**	**1**	**7**	**11**	**1**	**7**	**4**	**18**	**1**	**16**	
machine	setup info	setup info	setup info	setup info	setup info	setup info	setup info	setup info	setup info	setup info	
architecture	amd64	amd64	amd64	amd64	amd64	amd64	amd64	amd64	amd64	amd64	
host	ot73	ot72	ot92	ot48	ot71	ot91	ot62	ot41	ot93	ot42	
openbsd-linux-tcp4bench-parallel	PASS	PASS	XPASS	PASS	XPASS	XPASS	PASS	XPASS	PASS	XPASS	tcpbench -S1000000 -t10 -n100 10.10.22.4
linux-openbsd-linux-tcp4iperf-parallel-copy	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	PASS	PASS	PASS	ssh root@lt40 iperf3 -i1 -c10.10.21.2 -w200k -P15 -t10
linux-openbsd-linux-tcp4iperf-parallel-forward	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	PASS	PASS	PASS	ssh root@lt40 iperf3 -c10.10.22.4 -w200k -P15 -t10
linux-openbsd-linux-tcp4iperf-parallel-splice	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	PASS	PASS	PASS	ssh root@lt40 iperf3 -c10.10.21.2 -w200k -P15 -t10
linux-openbsd-linux-tcp6iperf-parallel-copy	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	PASS	PASS	PASS	ssh root@lt40 iperf3 -i1 -6 -cfdd7:e83e:66bd:1021::2 -w200k -P15 -t10
linux-openbsd-linux-tcp6iperf-parallel-forward	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	PASS	PASS	PASS	ssh root@lt40 iperf3 -6 -cfdd7:e83e:66bd:1022::4 -w200k -P15 -t10
linux-openbsd-linux-tcp6iperf-parallel-splice	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	PASS	PASS	PASS	ssh root@lt40 iperf3 -6 -cfdd7:e83e:66bd:1021::2 -w200k -P15 -t10
openbsd-linux-tcp4iperf-parallel-send	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	PASS	PASS	PASS	iperf3 -c10.10.22.4 -w200k -P15 -t10
openbsd-linux-tcp6iperf-parallel-send	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	PASS	PASS	PASS	iperf3 -6 -cfdd7:e83e:66bd:1022::4 -w200k -P15 -t10
openbsd-linux-tcp6bench-parallel	PASS	PASS	PASS	PASS	PASS	PASS	PASS	XPASS	PASS	XPASS	tcpbench -S1000000 -t10 -n100 fdd7:e83e:66bd:1022::4
linux-openbsd-linux-tcp6bench-single	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	XPASS	PASS	ssh root@lt40 tcpbench -S1000000 -t10 fdd7:e83e:66bd:1022::4
openbsd-linux-tcp4bench-single	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	XPASS	tcpbench -S1000000 -t10 10.10.22.4
linux-openbsd-linux-udp4bench-single-frag	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -f2 -croot@lt40 -sroot@lt43 -a10.10.22.40 -t10
linux-openbsd-linux-udp6bench-single-frag	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -f2 -croot@lt40 -sroot@lt43 -afdd7:e83e:66bd:1
linux-openbsd-udp4bench-single-frag	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -f2 -croot@lt40 -a10.10.21.20 -t10 udpbench
linux-openbsd-udp6bench-single-frag	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -f2 -croot@lt40 -afdd7:e83e:66bd:1021::20 -t10
linux-openbsd-linux-tcp4bench-parallel	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	XPASS	ssh root@lt40 tcpbench -S1000000 -t10 -n100 10.10.22.4
linux-openbsd-linux-tcp4bench-single	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	ssh root@lt40 tcpbench -S1000000 -t10 10.10.22.4
linux-openbsd-linux-tcp4copy-parallel	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -N10 -croot@lt40 -sroot@lt43 -A10.10.21.20 -a1
linux-openbsd-linux-tcp4copy-single	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -croot@lt40 -sroot@lt43 -A10.10.21.20 -a10.10.
linux-openbsd-linux-tcp4splice-parallel	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -N10 -croot@lt40 -sroot@lt43 -A10.10.21.20 -a1
linux-openbsd-linux-tcp4splice-single	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -croot@lt40 -sroot@lt43 -A10.10.21.20 -a10.10.
linux-openbsd-linux-tcp6bench-parallel	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	XPASS	ssh root@lt40 tcpbench -S1000000 -t10 -n100 fdd7:e83e:66bd:1022::4
linux-openbsd-linux-tcp6copy-parallel	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -N10 -croot@lt40 -sroot@lt43 -Afdd7:e83e:66bd:
linux-openbsd-linux-tcp6copy-single	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -croot@lt40 -sroot@lt43 -Afdd7:e83e:66bd:1021:
linux-openbsd-linux-tcp6splice-parallel	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -N10 -croot@lt40 -sroot@lt43 -Afdd7:e83e:66bd:
linux-openbsd-linux-tcp6splice-single	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -b1000000 -croot@lt40 -sroot@lt43 -Afdd7:e83e:66bd:1021:
linux-openbsd-linux-udp4bench-parallel-empty	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f0 -i0 -N10 -croot@lt40 -sro
linux-openbsd-linux-udp4bench-parallel-frag	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f2 -i0 -N10 -croot@lt40 -sro
linux-openbsd-linux-udp4bench-parallel-mtu	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f1 -i0 -N10 -croot@lt40 -sro
linux-openbsd-linux-udp4bench-single-empty	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	netbench.pl -v -b1000000 -f0 -croot@lt40 -sroot@lt43 -a10.10.22.40 -t10
linux-openbsd-linux-udp4bench-single-mtu	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	netbench.pl -v -b1000000 -f1 -croot@lt40 -sroot@lt43 -a10.10.22.40 -t10
linux-openbsd-linux-udp4splice-empty	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f0 -i0 -N10 -croot@lt40 -sro
linux-openbsd-linux-udp4splice-mtu	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f1 -i0 -N10 -croot@lt40 -sro
linux-openbsd-linux-udp6bench-parallel-empty	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f0 -i0 -N10 -croot@lt40 -sro
linux-openbsd-linux-udp6bench-parallel-frag	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f2 -i0 -N10 -croot@lt40 -sro
linux-openbsd-linux-udp6bench-parallel-mtu	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f1 -i0 -N10 -croot@lt40 -sro
linux-openbsd-linux-udp6bench-single-empty	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	netbench.pl -v -b1000000 -f0 -croot@lt40 -sroot@lt43 -afdd7:e83e:66bd:1
linux-openbsd-linux-udp6bench-single-mtu	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	netbench.pl -v -b1000000 -f1 -croot@lt40 -sroot@lt43 -afdd7:e83e:66bd:1
linux-openbsd-linux-udp6splice-empty	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f0 -i0 -N10 -croot@lt40 -sro
linux-openbsd-linux-udp6splice-mtu	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	NOEXIT	netbench.pl -v -B1000000000 -b1000000 -d1 -f1 -i0 -N10 -croot@lt40 -sro
linux-openbsd-mcast4bench-parallel-mtu	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	PASS	XPASS	netbench.pl -v -B1000000000 -b1000000 -d1 -f1 -i0 -N10 -R10.10.21.2 -S1

genua.

Test Matrix

- network hardware and driver
  bge bnxt em ice igc ix ixl re vio (vmx)
- modify network setup
  jumbo nolro nopf notso
- stack pseudo devices
  bridge carp gif gif6 gre veb vlan vxlan wg
- 93 test cases
  icmp tcp udp splice mcast iperf (trex)
- various platforms
  Intel AMD vmm-vmd KVM-qemu AMD-SEV (sparc64) (vmware)

Interface Throughput

					2025-09-03 iface-ice0	2025-09-03 iface-ice0	2025-09-03 iface-ice0	2025-09-03 iface-ice0	2025-09-03 iface-ice0	2025-09-03 iface-ice0	2025-09-03 iface-ice0	2025-09-0 iface-ice 2025-09-0
						000	001	002	003	004	btrace-kstack.0	
					100 Gbit	100 Gbit	100 Gbit	100 Gbit	100 Gbit	100 Gbit	100 Gbit	100 Gbit
linux-openbsd-linux-tcp6copy-parallel	IPv6 TCP	Copy		parallel-10	33.9%	31.0%	29.6%	31.2%	35.3%	34.4%	36.9%	35.8
linux-openbsd-linux-udp4splice-empty	IPv4 UDP	Splice	Empty	parallel-10	0.3%	0.2%	0.2%	0.1%	0.3%	0.3%	0.3%	0.2
linux-openbsd-linux-udp6splice-empty	IPv6 UDP	Splice	Empty	parallel-10	0.1%	0.3%	0.1%	0.3%	0.3%	0.3%	0.2%	0.3
linux-openbsd-linux-udp4splice-mtu	IPv4 UDP	Splice	MTU	parallel-10	7.4%	7.6%	7.2%	6.8%	7.3%	7.0%	7.0%	7.8
linux-openbsd-linux-udp6splice-mtu	IPv6 UDP	Splice	MTU	parallel-10	6.9%	7.8%	7.7%	7.7%	7.3%	7.4%	7.5%	7.5
linux-openbsd-mcast4bench-parallel-mtu	IPv4 MCast	Receive	MTU	parallel-10	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0
linux-openbsd-mcast6bench-parallel-mtu	IPv6 MCast	Receive	MTU	parallel-10	9.3%	9.7%	9.8%	9.9%	9.9%	10.0%	9.7%	8.8
openbsd-linux-mcast4bench-parallel-mtu	IPv4 MCast	Send	MTU	parallel-10	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0
openbsd-linux-mcast6bench-parallel-mtu	IPv6 MCast	Send	MTU	parallel-10	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0%	10.0
linux-openbsd-linux-tcp4iperf-parallel-forward	IPv4 TCP	Forward		parallel-15	47.9%	45.2%	48.6%	47.5%	46.4%	46.5%	45.9%	46.9
linux-openbsd-linux-tcp4iperf-parallel-splice	IPv4 TCP	Splice		parallel-15	24.5%	31.5%	31.9%	30.2%	28.2%	28.7%	30.9%	31.2
linux-openbsd-linux-tcp4iperf-parallel-copy	IPv4 TCP	Copy		parallel-15	48.7%	47.8%	48.3%	48.2%	47.7%	47.0%	48.4%	47.3
linux-openbsd-tcp4iperf-parallel-receive	IPv4 TCP	Receive		parallel-15	66.7%	62.4%	63.9%	61.3%	65.9%	64.1%	65.7%	62.3
openbsd-linux-tcp4iperf-parallel-send	IPv4 TCP	Send		parallel-15	48.6%	48.5%	48.6%	48.8%	48.6%	49.4%	50.0%	48.6
linux-openbsd-linux-tcp6iperf-parallel-forward	IPv6 TCP	Forward		parallel-15	44.7%	44.9%	45.7%	45.1%	44.9%	44.7%	45.5%	44.4
linux-openbsd-linux-tcp6iperf-parallel-splice	IPv6 TCP	Splice		parallel-15	31.0%	27.2%	30.8%	25.1%	26.0%	26.6%	26.6%	28.3
linux-openbsd-linux-tcp6iperf-parallel-copy	IPv6 TCP	Copy		parallel-15	43.9%	44.2%	45.4%	44.5%	44.6%	45.7%	43.6%	45.5
linux-openbsd-tcp6iperf-parallel-receive	IPv6 TCP	Receive		parallel-15	63.4%	63.9%	64.2%	65.5%	65.9%	62.7%	61.6%	62.4
openbsd-linux-tcp6iperf-parallel-send	IPv6 TCP	Send		parallel-15	48.5%	47.5%	47.0%	47.6%	47.3%	47.5%	47.2%	47.5
linux-openbsd-linux-tcp4iperf-multiple-forward	IPv4 TCP	Forward		parallel-30	46.5%	47.7%	46.9%	46.7%	47.6%	47.1%	9.4%	46.1
linux-openbsd-linux-tcp4iperf-multiple-splice	IPv4 TCP	Splice		parallel-30	9.4%	20.3%	19.7%	19.5%	20.5%	19.2%	9.4%	20.1
linux-openbsd-linux-tcp4iperf-multiple-copy	IPv4 TCP	Copy		parallel-30	41.1%	42.9%	41.4%	44.6%	44.4%	42.3%	9.4%	44.0
linux-openbsd-tcp4iperf-multiple-receive	IPv4 TCP	Receive		parallel-30	60.7%	60.0%	61.0%	64.3%	62.0%	59.1%	59.6%	65.6
openbsd-linux-tcp4iperf-multiple-send	IPv4 TCP	Send		parallel-30	47.9%	47.2%	47.0%	47.9%	46.6%	46.8%	47.0%	47.3
linux-openbsd-linux-tcp6iperf-multiple-forward	IPv6 TCP	Forward		parallel-30	43.4%	9.3%	45.0%	44.8%	44.3%	45.9%	43.6%	9.3
linux-openbsd-linux-tcp6iperf-multiple-splice	IPv6 TCP	Splice		parallel-30	19.2%	20.4%	20.5%	19.6%	21.3%	19.9%	9.3%	
linux-openbsd-linux-tcp6iperf-multiple-copy	IPv6 TCP	Copy		parallel-30	9.3%	38.5%	9.3%	37.6%	39.2%	41.2%	35.1%	39.6
linux-openbsd-tcp6iperf-multiple-receive	IPv6 TCP	Receive		parallel-30	59.6%	54.8%	59.1%	52.6%	57.8%	59.7%	9.3%	58.9
openbsd-linux-tcp6iperf-multiple-send	IPv6 TCP	Send		parallel-30	45.5%	46.5%	46.5%	46.0%	46.1%	46.4%	46.0%	46.2
trex-tcp4-splice	IPv4 TCP	Splice		trex	0.7%	0.7%	0.7%	0.7%	0.7%	0.7%		0.7

PASS netlink test passed
FAIL netlink test failed to produce value
XPASS netlink test passed, error in netstat output
NOEXIT netlink test did not exit with code 0, test failed
NOTERM test did not terminate, aborted after timeout
NORUN test did not run, execute failed
NOLOG create log file for test output failed

genua.

TCP Segmentation Offload, TSO

- Large Receive Offload, LRO
  - hardware support
    ix vio vmx
  - software implementation, easy to expand
    ice ixl
- TCP Send Offload, TSO
  - hardware support
    bnxt em iavf ice igc ix ixl vio vmx

Socket Splicing

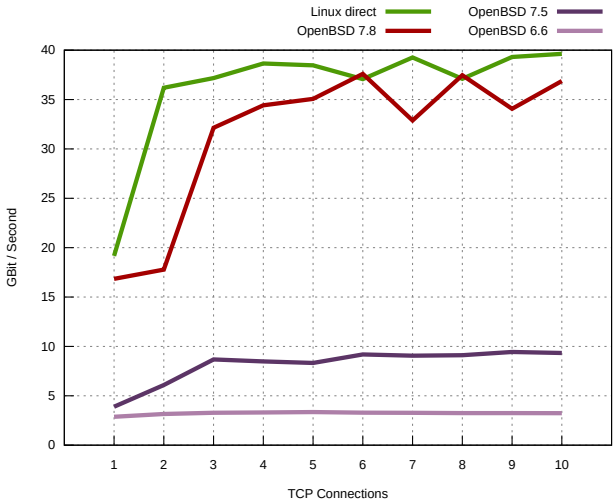

OpenBSD	release
7.8	Oct 2025
7.5	Apr 2024
6.6	Oct 2019

genua.

Know Your Optimization Goal

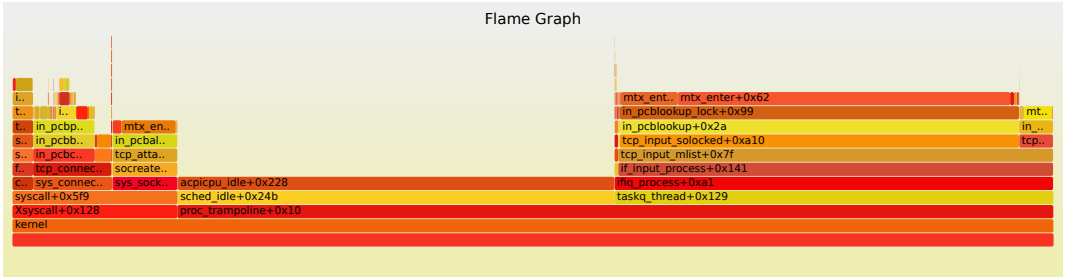

- Cisco TRex stateful http simple
- relayd(8) socket splicing
- 700 MBit per second throughput

genua.

Links

- http://bluhm.genua.de/test.html
- http://bluhm.genua.de/perform/results/perform.html
- http://bluhm.genua.de/netlink/results/latest.html
- https://github.com/bluhm/talk-protomp