

Measuring Performance on OpenBSD

Alexander Bluhm

bluhm@openbsd.org

EuroBSDCon, September 2019

Agenda

- 1 What did exist before?
- 2 How does it work?
- 3 What are the findings?
- 4 What is the Conclusion?

What did exist before?



How does it work?



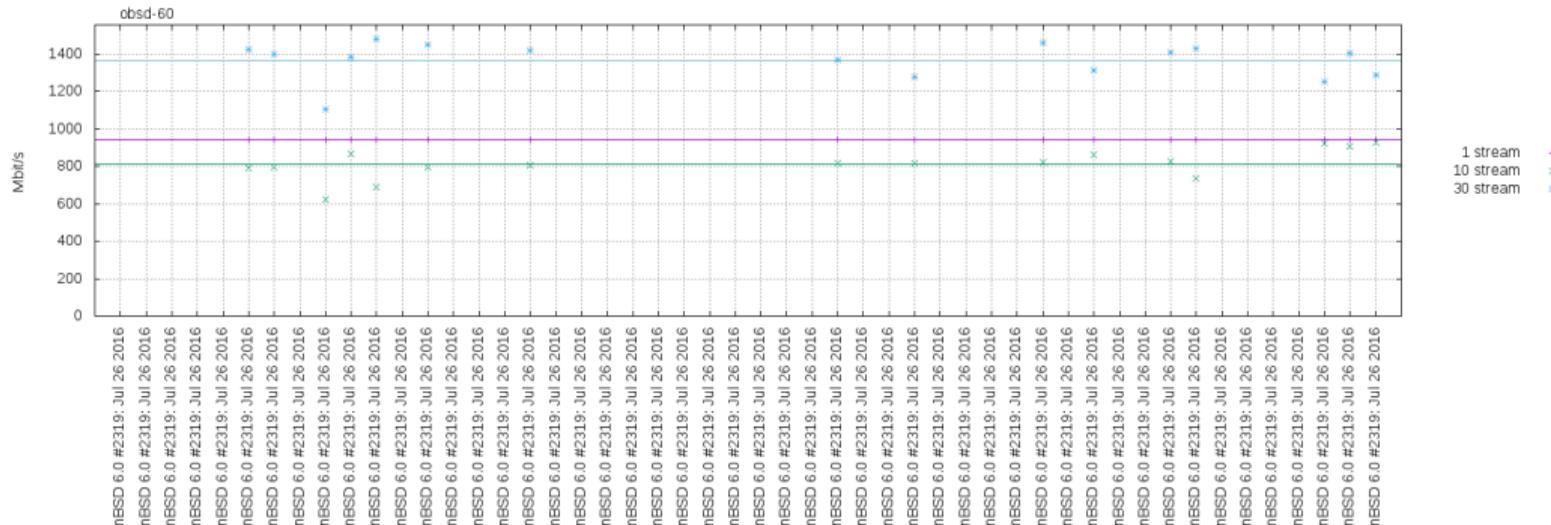
What are the findings?



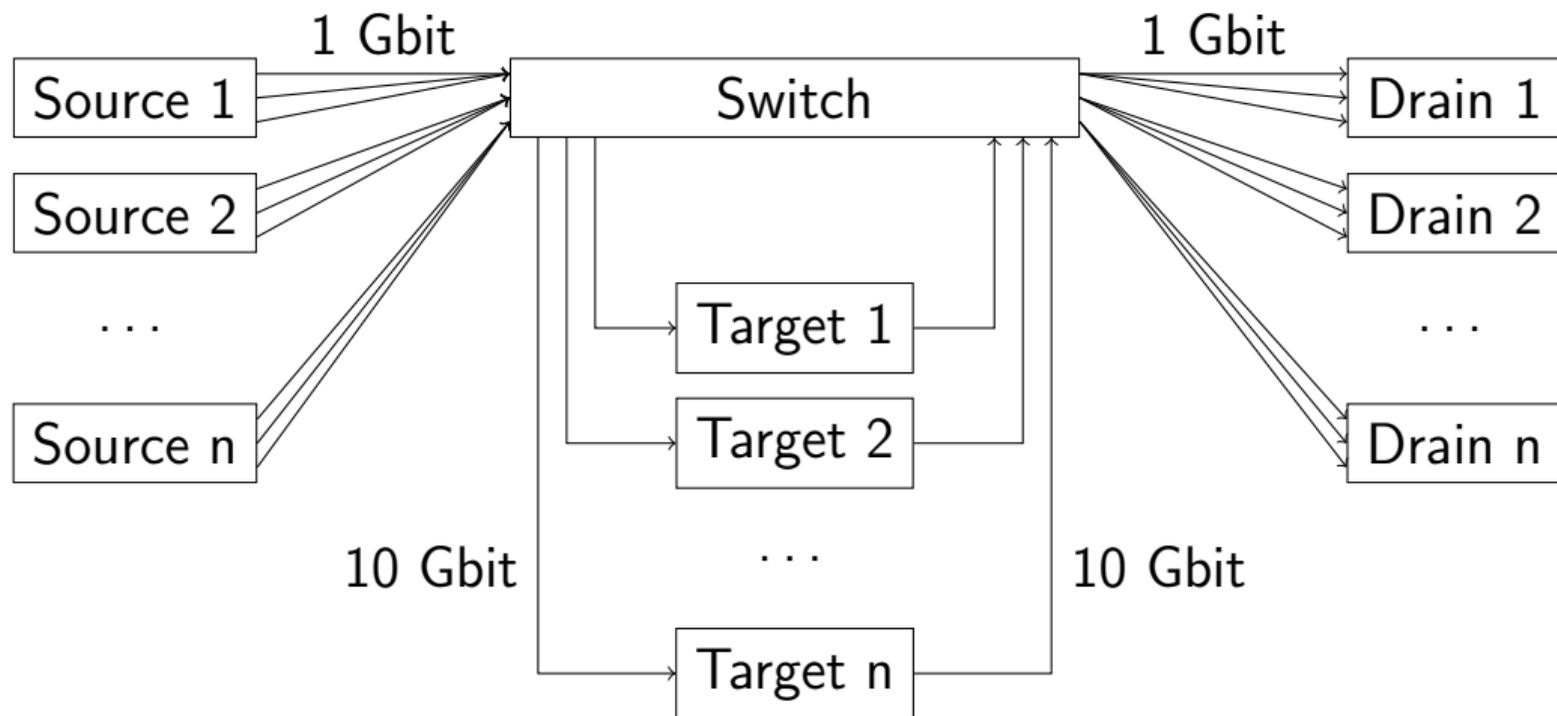
What is the Conclusion?



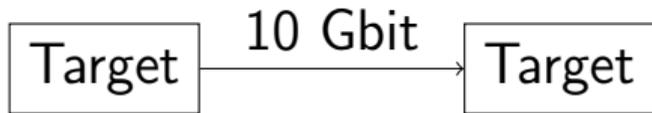
genua Firewall Testbed HPF



Multi User, Multi Purpose Hardware Setup



Performance Hardware Design



What did exist before?

○○○○●○

How does it work?

○○○○○○○○○○

What are the findings?

○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○

What is the Conclusion?

○○○○○

Existing Regression Tests

OpenBSD regress all test results

created at 2019-04-17T21:03:24Z

test [run info](#)

pass rate	96%	91%	96%	91%	98%	96%	90%	98%	96%	90%	98%	96%	91%	98%	97%
run at date	2019-04-17	2019-04-17	2019-04-17	2019-04-16	2019-04-16	2019-04-16	2019-04-15	2019-04-15	2019-04-15	2019-04-14	2019-04-14	2019-04-13	2019-04-13	2019-04-13	2019-04-13
machine build	snapshot														
architecture	arm64	armv7	amd64	armv7	amd64	i386	armv7	amd64	i386	armv7	amd64	arm64	armv7	amd64	i386
misc/posixtestsuite	NOTERM														
lib/libc	NOTERM	FAIL	PASS	FAIL	PASS	PASS	FAIL	PASS	PASS	FAIL	PASS	NOTERM	FAIL	PASS	PASS
lib/libm	FAIL	FAIL	PASS	FAIL	PASS	PASS	FAIL	PASS	PASS	FAIL	PASS	FAIL	FAIL	PASS	PASS
usr.sbin/ospfd	FAIL	FAIL	PASS	FAIL	PASS	PASS	FAIL	PASS	PASS	FAIL	PASS	FAIL	FAIL	PASS	PASS
lib/libcrypto	SKIP	NOEXIT	FAIL	NOEXIT	FAIL	FAIL	NOEXIT	FAIL	FAIL	NOEXIT	FAIL	SKIP	NOEXIT	PASS	PASS
sys/net/pflow	FAIL	PASS	NOTERM	PASS	FAIL	PASS	PASS	PASS							
gnu/egcs/gcc-bounds	SKIP	FAIL	SKIP	FAIL	FAIL	FAIL									
gnu/egcs/gcc-builtins	SKIP	FAIL	SKIP	FAIL	FAIL	FAIL									
lib/libpthread	FAIL	PASS	FAIL	PASS	PASS	PASS									
usr.bin/bc	FAIL	PASS	FAIL	PASS	PASS	PASS									
usr.bin/ssh	PASS	NOTERM	PASS	NOTERM	PASS	PASS									
usr.bin/ctfdump	SKIP	FAIL	PASS	FAIL	PASS	PASS	FAIL	PASS	PASS	FAIL	PASS	SKIP	FAIL	PASS	PASS
lib/libssl	PASS	FAIL	PASS	NOTERM	PASS	PASS	FAIL	PASS	PASS	NOTERM	PASS	PASS	FAIL	PASS	PASS
usr.sbin/syslogd	SKIP	SKIP	FAIL	SKIP	PASS	FAIL	SKIP	PASS	PASS	SKIP	PASS	SKIP	SKIP	PASS	PASS
sbin/disklabel	PASS	FAIL	PASS	FAIL	PASS	PASS									
sbin/pfctl	PASS	FAIL	PASS	FAIL	PASS	PASS									
sys/kern/ptrace	PASS	FAIL	PASS	FAIL	PASS	PASS									
sys/uvm/mmap_write_self	PASS	FAIL	PASS	FAIL	PASS	PASS									
usr.bin/mail	PASS	FAIL	PASS	FAIL	PASS	PASS									
usr.bin/openssl	PASS	FAIL	PASS	FAIL	PASS	PASS									
usr.sbin/rebound	PASS	FAIL	PASS	FAIL	PASS	PASS									
usr.sbin/snmpd	PASS	FAIL	PASS	FAIL	PASS	PASS									

Regression History for i386

history



OpenBSD regress ot1 test results

created at 2019-04-18T05:40:39Z

test run.info

pass rate	97%	96%	96%	97%	98%	97%	98%	98%	98%	98%
run at date	2019-04-18	2019-04-16	2019-04-15	2019-04-13	2019-04-11	2019-04-10	2019-04-09	2019-04-08	2019-04-06	2019-04-04
machine build	snapshot	custom								
architecture	i386									
misc/posixtestsuite	NOTERM									
lib/libcrypto	FAIL	FAIL	FAIL	PASS						
gnu/egcs/gcc-bounds	FAIL									
gnu/egcs/gcc-builtins	FAIL									
sys/net/pf_divert	FAIL	PASS	FAIL	FAIL	PASS	FAIL	PASS	PASS	PASS	PASS
sys/netinet/frag	FAIL	PASS	FAIL	FAIL	PASS	FAIL	PASS	PASS	PASS	PASS
lib/libc	PASS									
lib/libm	PASS									
usr.sbin/ospfd	PASS									
sys/net/pflow	PASS									
lib/libpthread	PASS									
usr.bin/bc	PASS									
sys/kern/flock	SKIP									
bin/ed	SKIP									
lib/libsndio	SKIP									
sys/arch/amd64	SKIP									
sys/arch/hppa	SKIP									
sys/arch/m88k	SKIP									
sys/arch/sparc64	SKIP									
sys/dev/kcov	SKIP									

severity



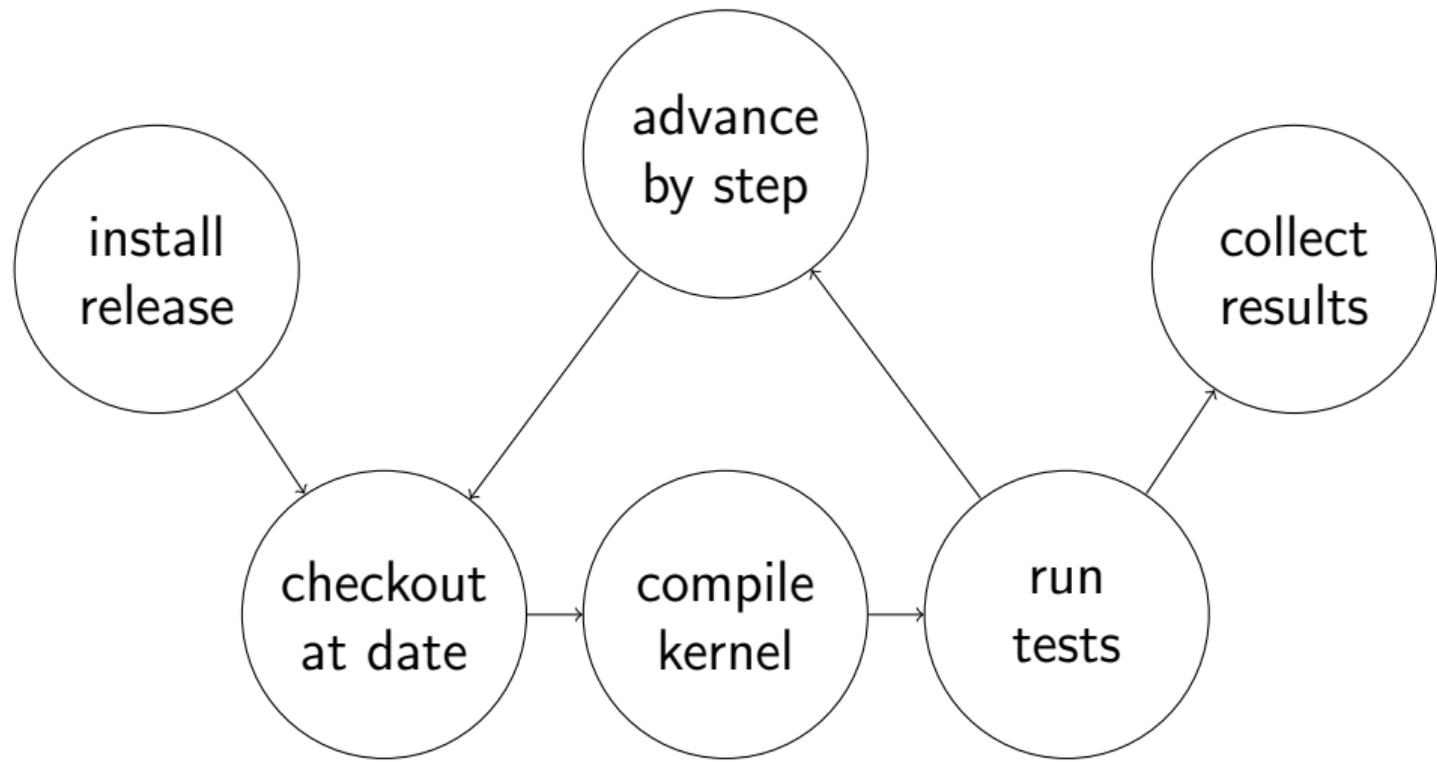
Agenda

- 1 What did exist before?
- 2 How does it work?**
- 3 What are the findings?
- 4 What is the Conclusion?

Performance Goals

- history
- reproducible
- details available
- drill down
- automatic

Performance History



Performance Tests Overview

run history

run detail

test command

run date

install version

cvcs checkout steps

OpenBSD perform all test results

```

created at 2019-04-26T02:37:12Z
test run info
run
host cores
release setup
first cvs checkout
last cvs checkout
steps
repetitions
iperf3 -c10.3.0.33 -w1m -t10
iperf3 -c10.3.0.33 -w1m -t10 -R
tcpbench -S1000000 -t10 10.3.0.33
tcpbench -S1000000 -t10 -n100 10.3.0.33
iperf3 -c10.3.0.33 -u -b10G -w1m -t10
iperf3 -c10.3.0.33 -u -b10G -w1m -t10 -R
iperf3 -c10.3.2.35 -w1m -t10
iperf3 -c10.3.2.35 -w1m -t10 -R
tcpbench -S1000000 -t10 10.3.2.35
tcpbench -S1000000 -t10 -n100 10.3.2.35
iperf3 -c10.3.2.35 -u -b0 -w1m -t10
iperf3 -c10.3.2.35 -u -b0 -w1m -t10 -R
iperf3 -c10.3.2.35 -u -b10G -w1m -t10
iperf3 -c10.3.2.35 -u -b10G -w1m -t10 -R
time -lp make -CGENERIC.MP -j4 -s
time -lp fs_mark -dfs_mark -DB -N16 -n256 -t8

```

	2019-04-25	2019-04-24	2019-04-20	2019-04-18	2019-04-17	2019-03-30	2019-03-23	2019-03-20	2019-03-20
ot12/8	ot14/4	ot14/4	ot12/8	ot14/4	ot14/4	ot14/4	ot14/4	ot14/4	ot14/4
6.5/install	6.4/install								
2019-04-16	2018-10-11	2019-04-13	2018-10-11	2019-03-30	2019-03-23	2019-03-16	2018-10-11	2018-10-11	2018-10-11
2019-04-17	2018-11-16	2019-04-20	2018-11-13	2019-04-13	2019-03-30	2019-03-23	2018-10-15	2018-10-15	2018-10-15
9/3hour	37/1day	8/1day	34/1day	15/1day	8/1day	8/1day	5/1day	5/1day	5/1day
5/reorder	5/reorder	5/reorder	5/keep	5/reorder	5/reorder	5/reorder	10/reorder	10/reboot	
iperf3 -c10.3.0.33 -w1m -t10	NOEXIT		PASS						
iperf3 -c10.3.0.33 -w1m -t10 -R	NOEXIT		PASS						
tcpbench -S1000000 -t10 10.3.0.33	FAIL		PASS						
tcpbench -S1000000 -t10 -n100 10.3.0.33	FAIL		PASS						
iperf3 -c10.3.0.33 -u -b10G -w1m -t10	NOEXIT		PASS						
iperf3 -c10.3.0.33 -u -b10G -w1m -t10 -R	NOEXIT		PASS						
iperf3 -c10.3.2.35 -w1m -t10	PASS	PASS		PASS	PASS	PASS	PASS	PASS	PASS
iperf3 -c10.3.2.35 -w1m -t10 -R	PASS	PASS		PASS	PASS	PASS	PASS	PASS	PASS
tcpbench -S1000000 -t10 10.3.2.35	PASS	PASS		PASS	PASS	PASS	PASS	PASS	PASS
tcpbench -S1000000 -t10 -n100 10.3.2.35	PASS	PASS		PASS	PASS	PASS	PASS	PASS	PASS
iperf3 -c10.3.2.35 -u -b0 -w1m -t10									
iperf3 -c10.3.2.35 -u -b0 -w1m -t10 -R									
iperf3 -c10.3.2.35 -u -b10G -w1m -t10	PASS	PASS		PASS	PASS	PASS	PASS	PASS	PASS
iperf3 -c10.3.2.35 -u -b10G -w1m -t10 -R	PASS	PASS		PASS	PASS	PASS	PASS	PASS	PASS
time -lp make -CGENERIC.MP -j4 -s									
time -lp fs_mark -dfs_mark -DB -N16 -n256 -t8	PASS		PASS						

Performance Run at Date

run and install log

OpenBSD perform 2019-02-04 test results

```

created at          2019-04-26T02:37:12Z
run at             2019-02-04T15:10:35Z
run log
test host with cpu cores ot12/8
machine release set 6.2/install info
steps              20/4usok
  
```

average numbers

```

cvs checkout
machine
kernel build
kernel commits
build quirks
repetitions
iperf3_c10.3.0.33_w1m_t10
sender
receiver
iperf3_c10.3.0.33_w1m_t10_R
sender
receiver
tcpbench_S1000000_t10_10.3.0.33
sender
tcpbench_S1000000_t10_n100_10.3.0.33
sender
iperf3_c10.3.0.33_u_b10G_w1m_t10
iperf3_c10.3.0.33_u_b10G_w1m_t10_R
time_ip_make_CGENERIC.MP_j8_s
  
```

unstable results

checkout date
kernel commits
build quirks

2017-10-04	2017-10-11	2017-10-18	2017-10-25	2017-11-01	2017-11-08	2017-11-15	2017-11-22	2017-11-29
build info								
version								
quirks/A,B,C								
5/reorder								
PASS								
4070000000	3990000000	4040000000	3990000000	3980000000	3930000000	4020000000	3930000000	3890000000
4070000000	3990000000	4040000000	3990000000	3980000000	3930000000	4020000000	3930000000	3890000000
PASS								
4090000000	4000000000	4090000000	4080000000	4040000000	4010000000	3970000000	3970000000	3990000000
4090000000	4000000000	4090000000	4080000000	4040000000	4010000000	3970000000	3970000000	3990000000
PASS								
4046336625	3988681500	3997288125	3945885875	3957848500	3873461875	3943817875	3842013250	3838678000
PASS								
4108316125	3980434125	4002707125	3985919875	3939937625	3923031875	3945349500	3863565375	3877594000
PASS								
PASS								
PASS								

Performance Repeat at CVS Checkout

reboot log

single result

repeat count

outlier

standard deviation

OpenBSD perform 2019-02-04 cvs 2017-10-18 test results

```

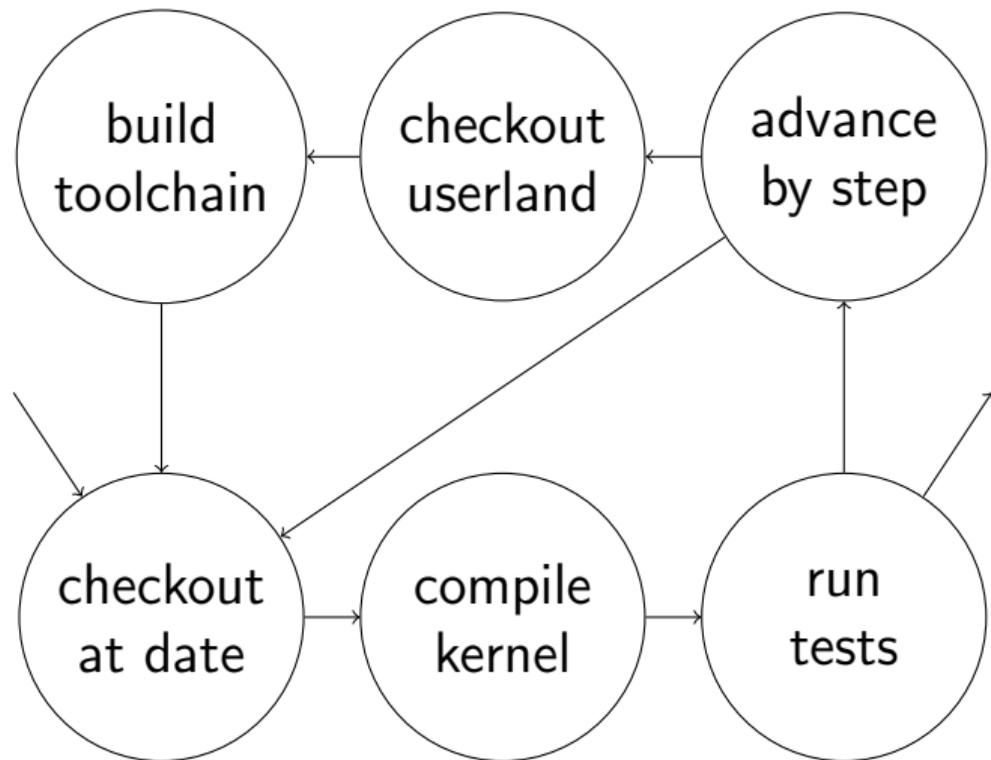
created at      2019-04-26T02:37:12Z
run at         2019-02-04T15:10:35Z
test host with cpu cores 0128
cvs checkout at 2017-10-18T00:00:00Z
repetitions    5
reorder info
repeat
machine
iperf3_c10.3.0.33_w1m_t10
sender
receiver
iperf3_c10.3.0.33_w1m_t10_R
sender
receiver
tcpbench_s1000000_t10_10.3.0.33
sender
receiver
tcpbench_s1000000_t10_n100_10.3.0.33
sender
iperf3_c10.3.0.33_u_b10G_w1m_t10
iperf3_c10.3.0.33_u_b10G_w1m_t10_R
time_lp_make_CGENERIC.MP_j8_s
real
user
sys
time_lp_fs_mark_dfs_mark_D8_N16_n256_t8
files
  
```

	000	001	002	003	004	unit	mean	minimum	maximum	deviation	relative
iperf3_c10.3.0.33_w1m_t10	PASS	PASS	PASS	PASS	PASS	bits/sec	3466000000	1600000000	4040000000	935427175.145131	0.26988666334250
iperf3_c10.3.0.33_w1m_t10_R	PASS	PASS	PASS	PASS	PASS	bits/sec	3492000000	1580000000	4090000000	961590349.36921	0.275369515856017
iperf3_c10.3.0.33_u_b10G_w1m_t10	PASS	PASS	PASS	PASS	PASS	bits/sec	3439756225	1602793875	3997288125	920681160.335085	0.267658689787429
iperf3_c10.3.0.33_u_b10G_w1m_t10_R	PASS	PASS	PASS	PASS	PASS	bits/sec	3448646150	1595932875	4002707125	927834572.54814	0.269043135254547
time_lp_make_CGENERIC.MP_j8_s	87.10	85.02	85.78	86.42	84.98	sec	85.66	84.98	87.10	0.817018971627953	0.0095157112931278
time_lp_fs_mark_dfs_mark_D8_N16_n256_t8	PASS	PASS	PASS	PASS	PASS	sec	177.12	172.14	180.00	2.63616084486513	0.0148899618190021
files	146.0	146.7	146.4	148.9	142.6	1/sec	146.12	142.6	148.9	2.02721483814617	0.013873630154299

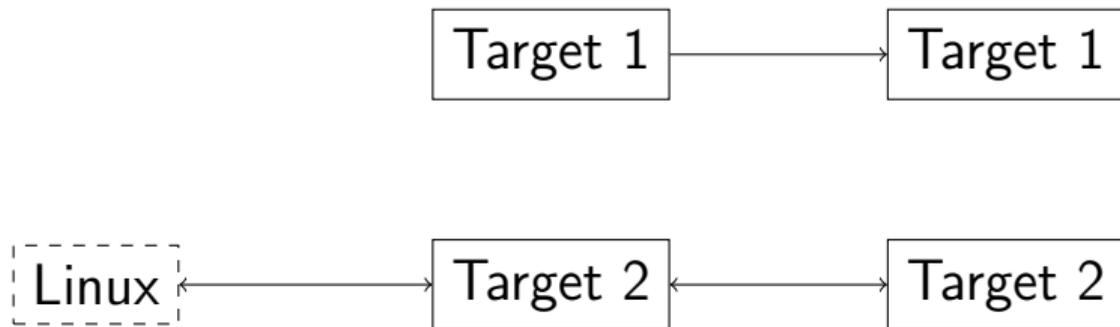
Quirks from 6.2 to 6.3

- A OpenBSD/amd64 6.2 release
- B fix cvs vendor branch checkout
- C clang update LLVM to 5.0.0
- D pfctl pf packet rate matching
- E move kernel source file dwiic.c
- F pfctl pf divert type
- G sysctl struct vfsconf
- H clang update LLVM to 5.0.1
- I pfctl pf syncookies
- J OpenBSD/amd64 6.3 release

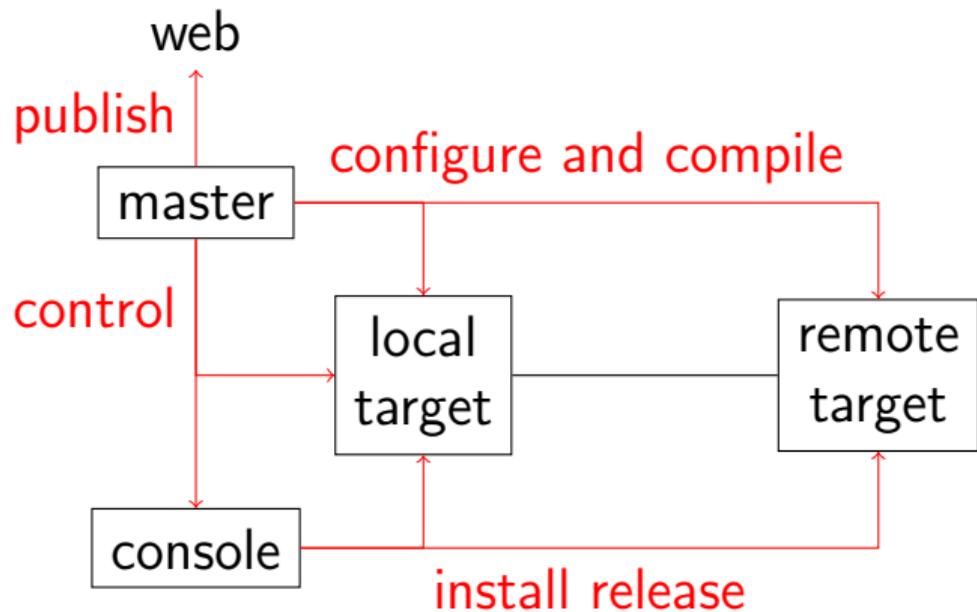
Build Quirks



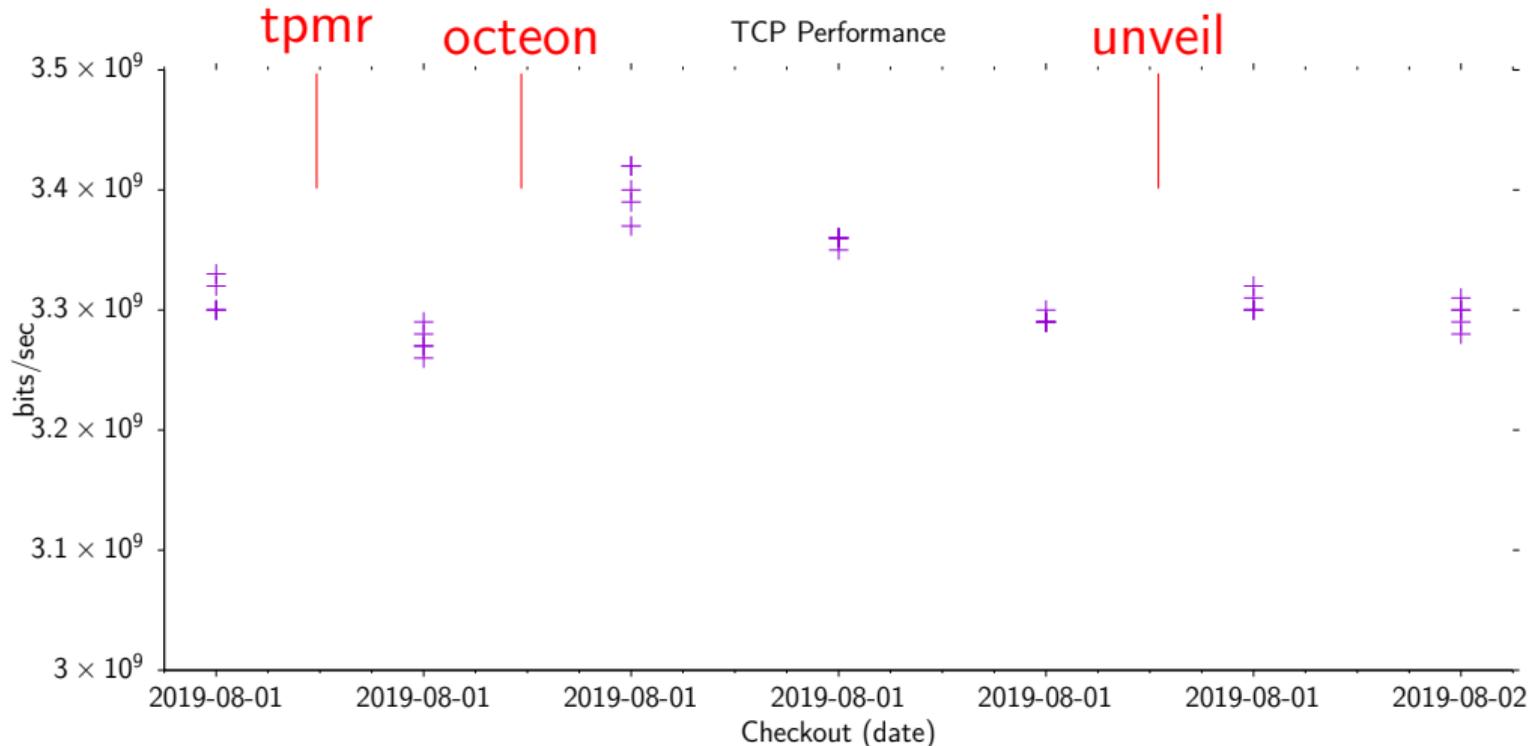
Performance Hardware



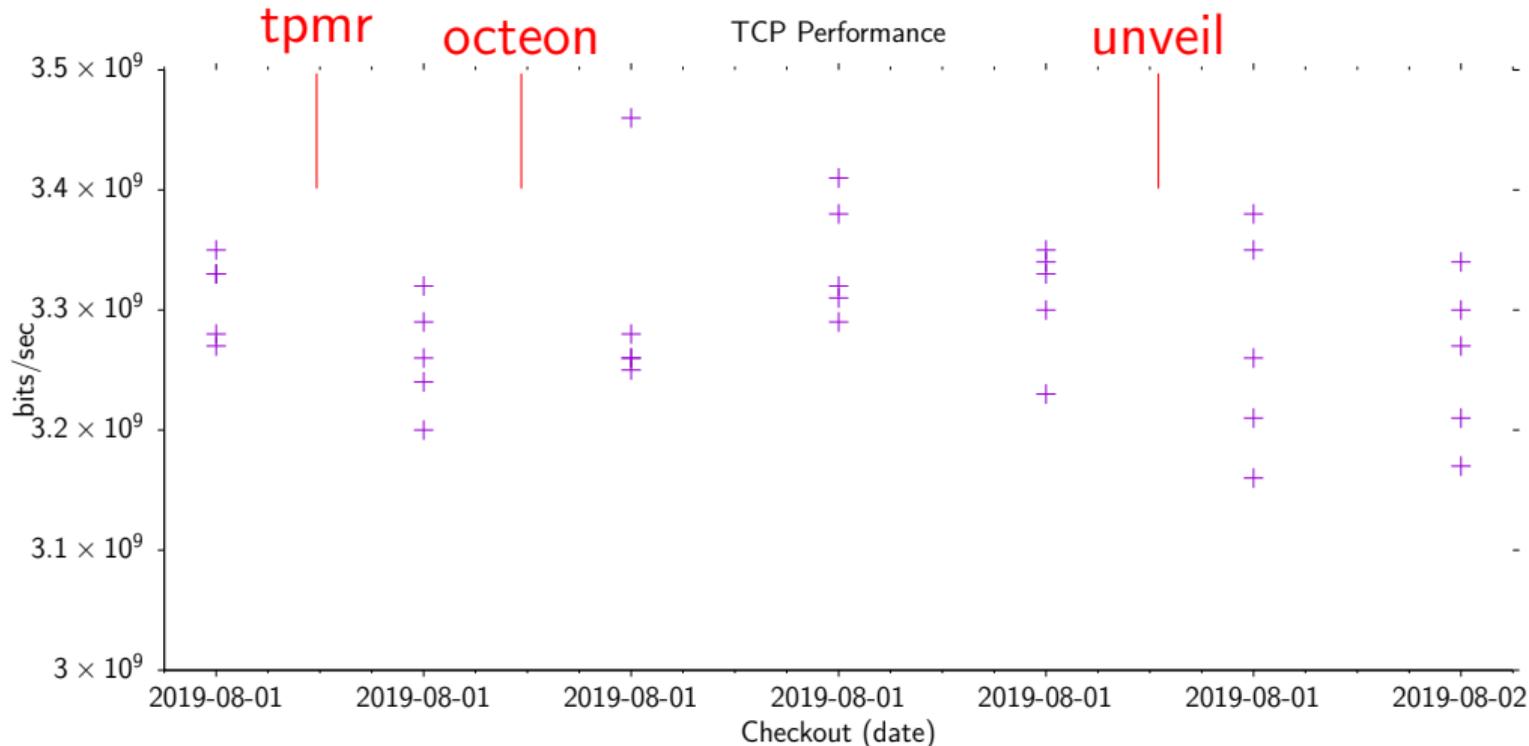
Performance Master



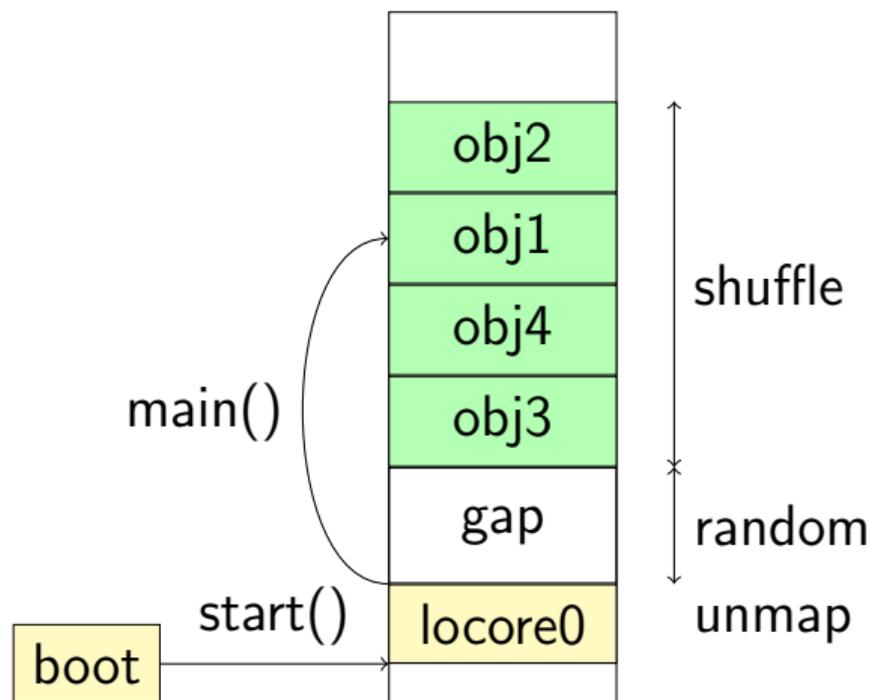
6.5, 1 Day, 5 Tests, Keep Machine Running



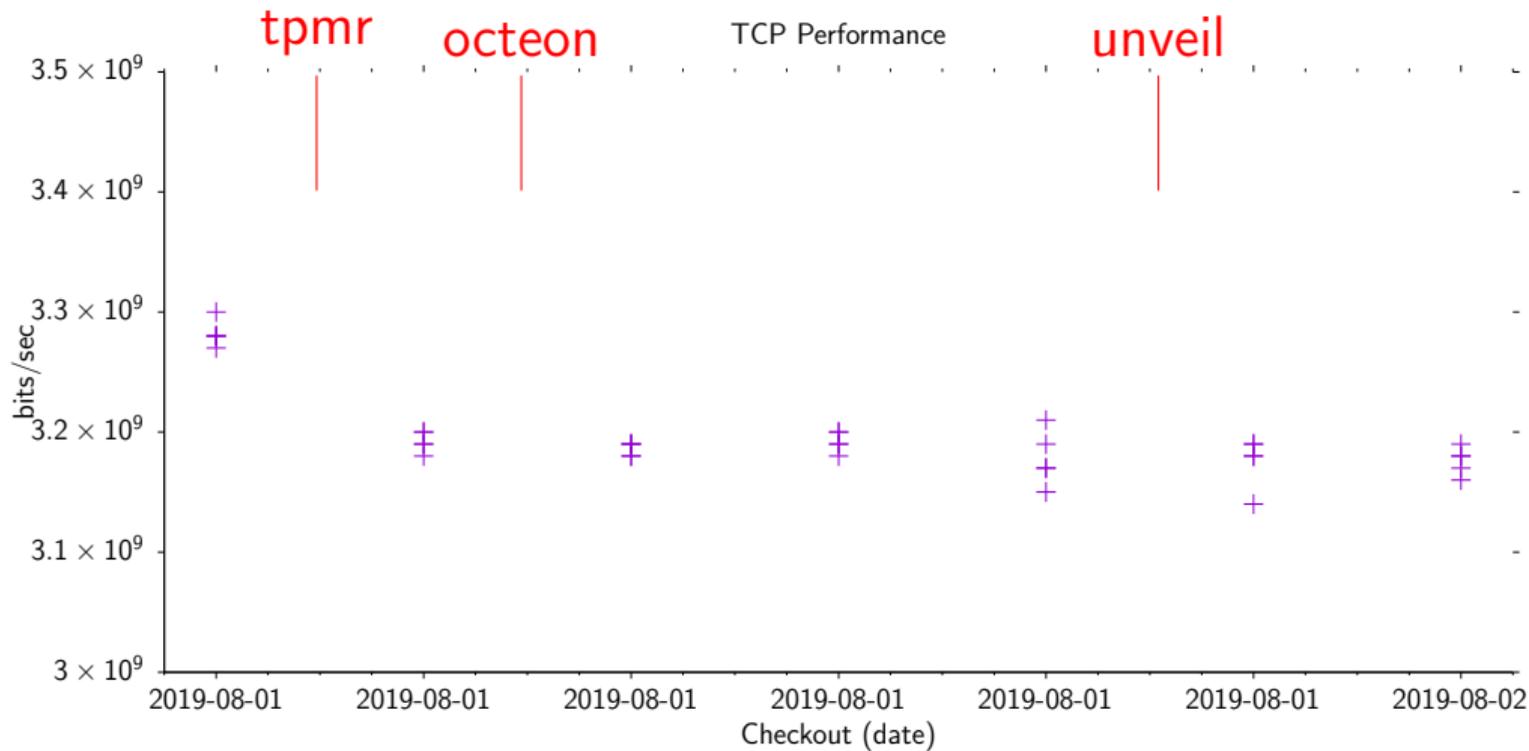
6.5, 1 Day, 5 Tests, Link and Reorder Kernel



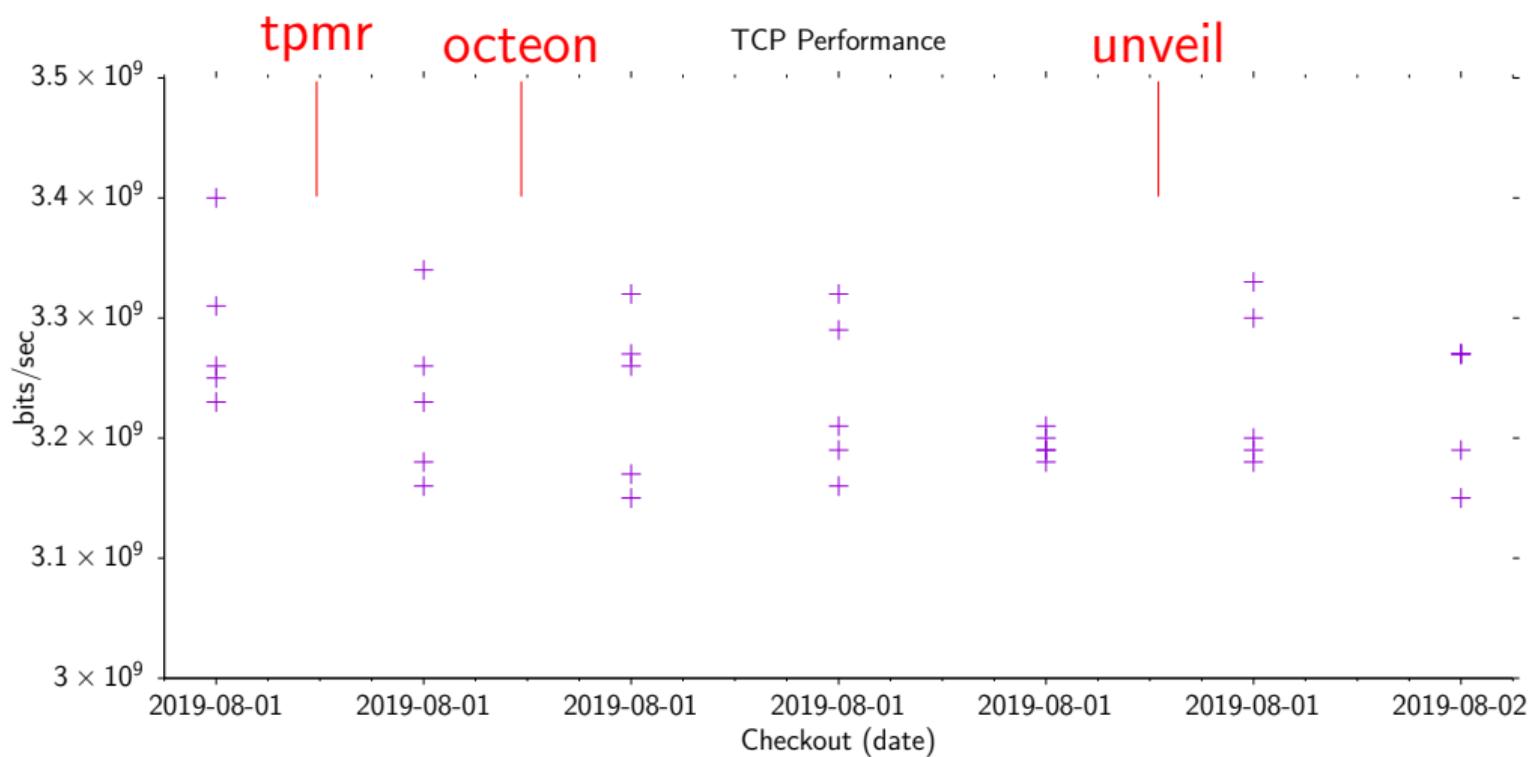
KARL Kernel Address Randomized Link



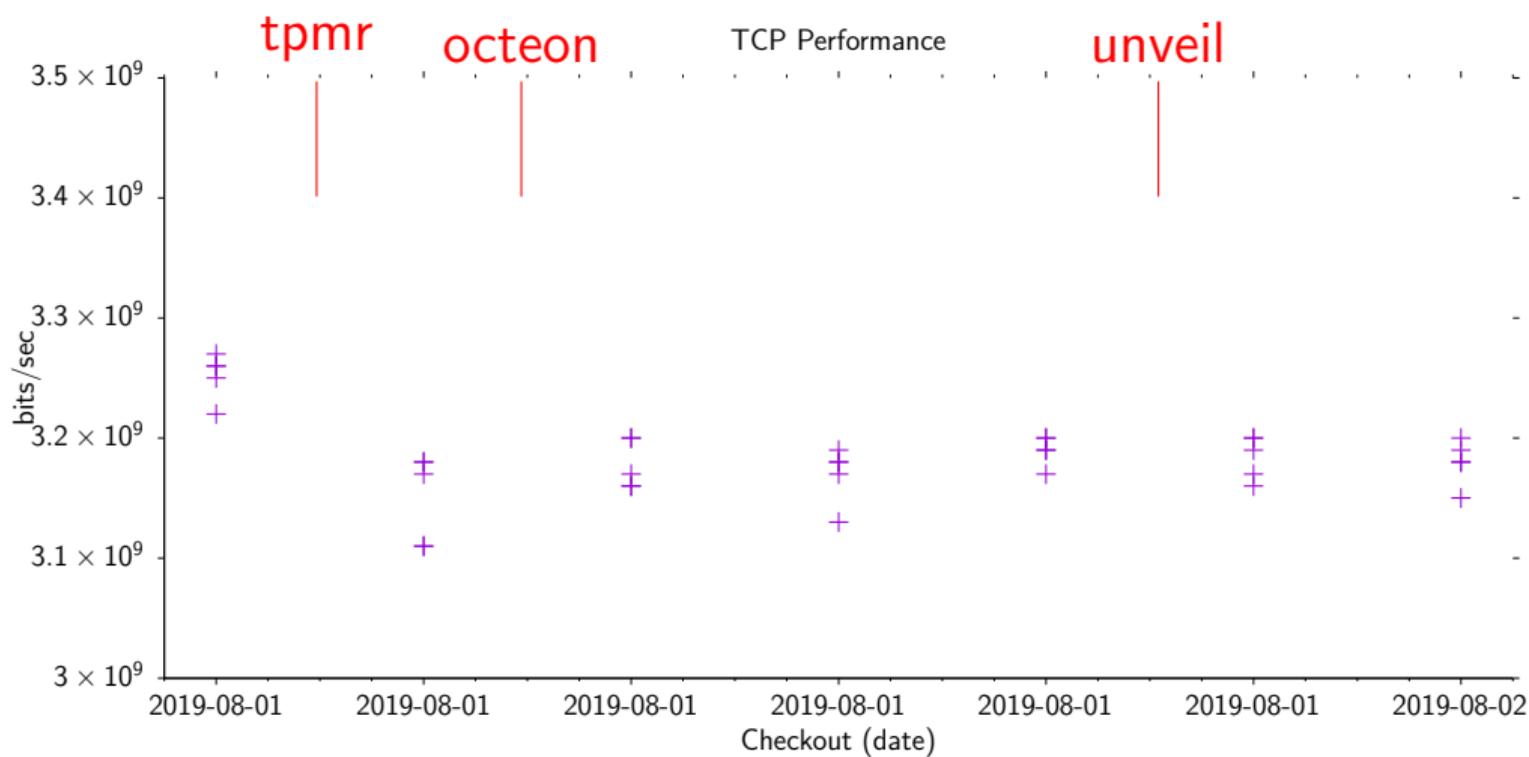
6.5, 1 Day, 5 Tests, Sort Objects, Fixed Gap



6.5, 1 Day, 5 Tests, Sort Objects, Random Gap



6.5, 1 Day, 5 Tests, Align Sorted Objects, Fixed Gap

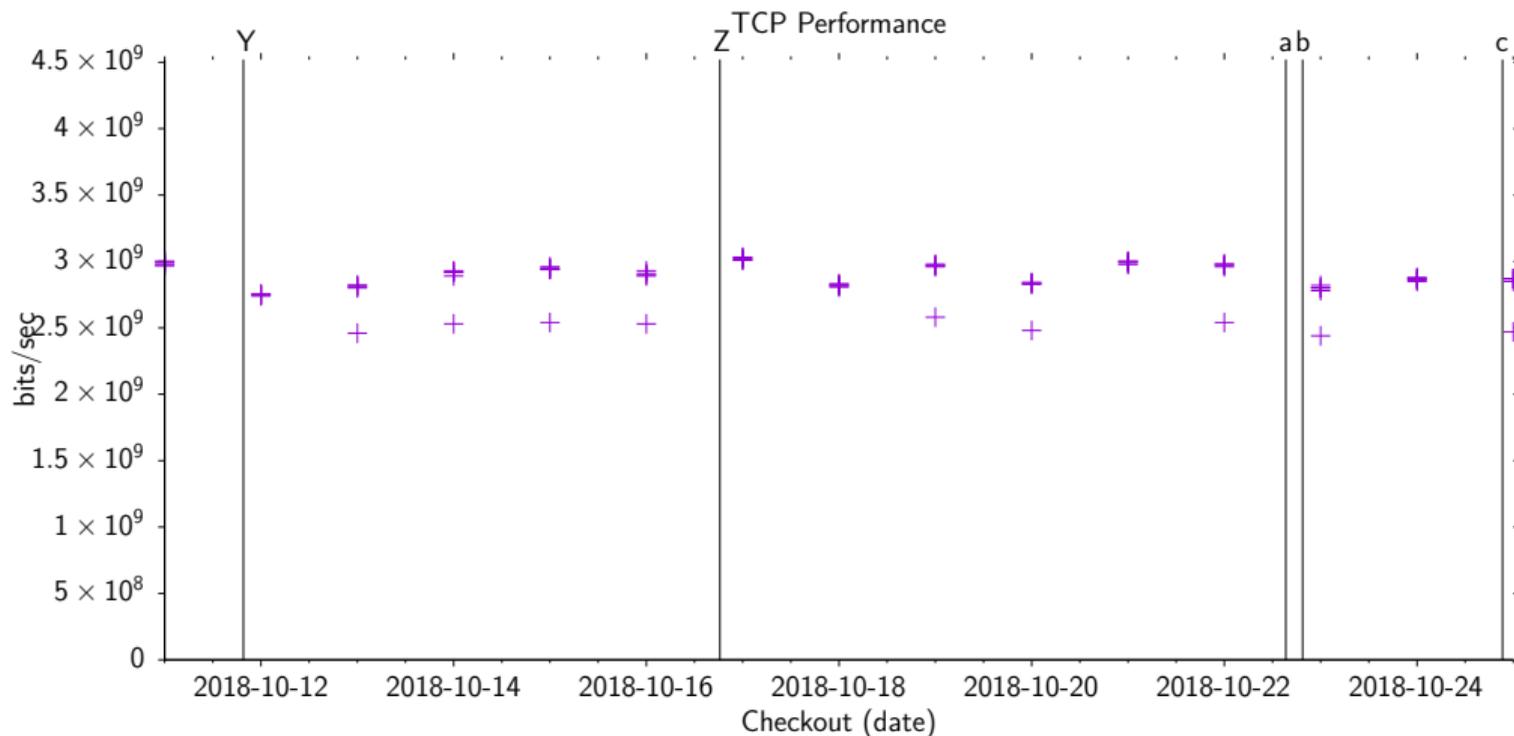


Kernel Symbol Table

nm bsd, diff, diffstat

	sort	align
unveil	+169	+13
	-169	-13
tpmr	+25997	+28731
	-25983	-28717

6.4, 15 Days, 5 Tests, 2 CPU Sockets, Keep running



2 CPU Sockets, Repeat, Keep running

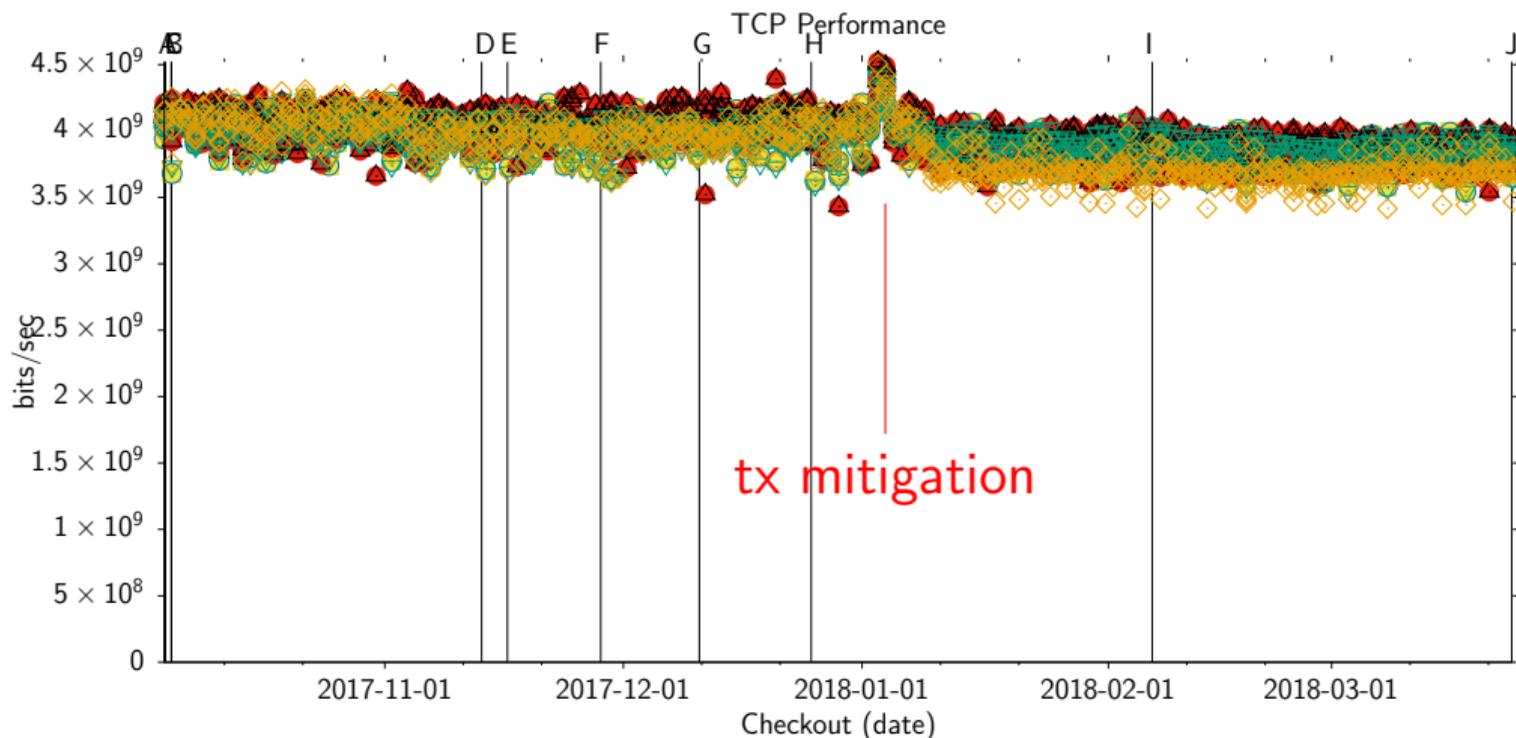
OpenBSD perform 2019-04-30 cvs 2018-10-13 test results

created at 2019-05-01T10:15:32Z
 run at 2019-04-30T19:11:10Z
 test host with cpu cores ot12/8
 cvs checkout at 2018-10-13T00:00:00Z
 repetitions [5/keep info](#)

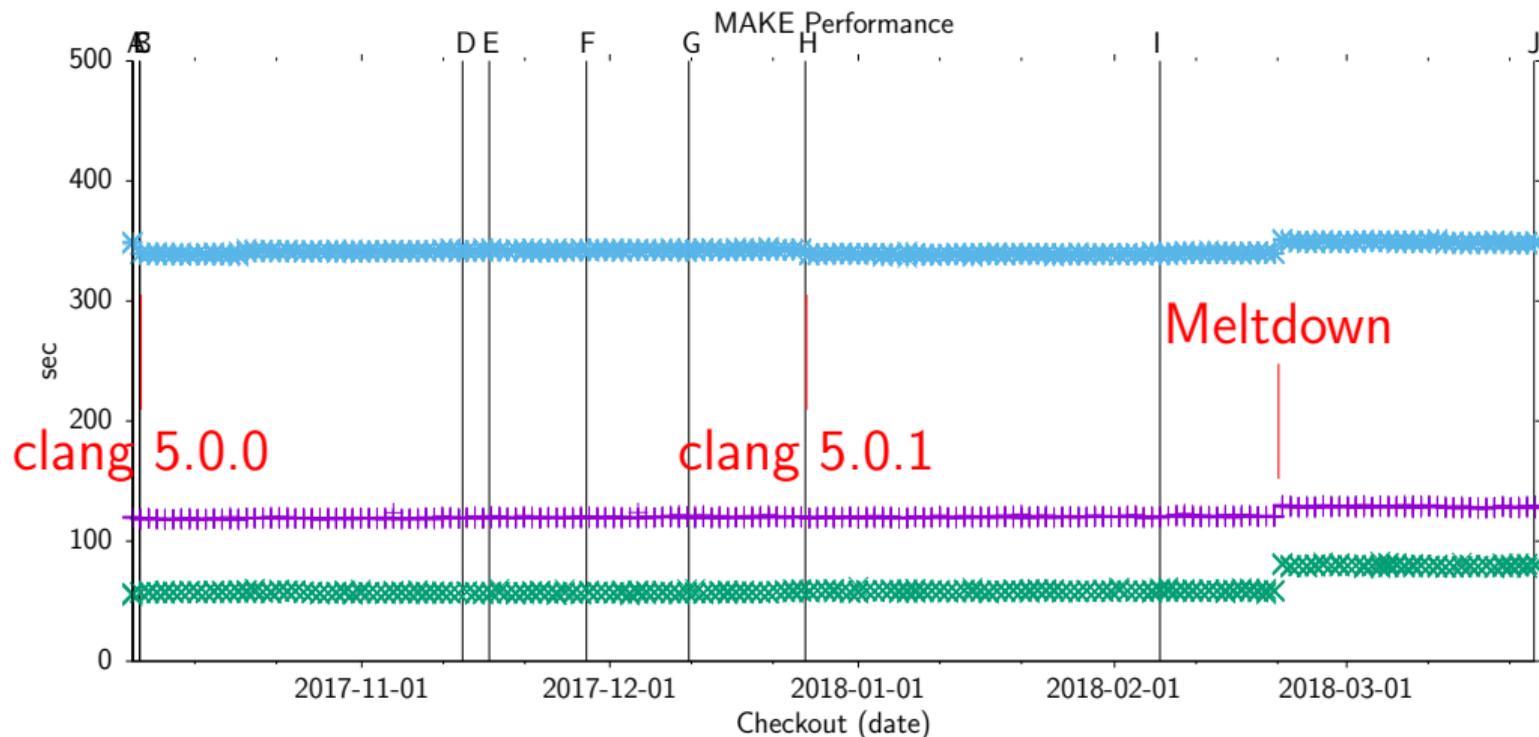
second cycle

repeat	000	001	002	003	004							
machine	keep	keep	keep	keep	keep	unit	mean	minimum	maximum	deviation	relative	
iperf3_c10.3.0.33_w1m_t10	PASS	PASS	PASS	PASS	PASS	bits/sec	2742000000	2460000000	2820000000	141194900.757782	0.0514933992552087	
sender	2800000000	2460000000	2820000000	2820000000	2810000000							
receiver	2750000000	2420000000	2770000000	2760000000	2760000000	bits/sec	2692000000	2420000000	2770000000	136146979.40094	0.0505746580241234	
iperf3_c10.3.0.33_w1m_t10_R	PASS	PASS	PASS	PASS	PASS	unit <th>mean</th> <th>minimum</th> <th>maximum</th> <th>deviation</th> <th>relative</th> <th></th>	mean	minimum	maximum	deviation	relative	
sender	2840000000	2800000000	2840000000	2840000000	2830000000	bits/sec	2830000000	2800000000	2840000000	15491933.3848297	0.00547418140806702	
receiver	2890000000	2850000000	2890000000	2890000000	2880000000	bits/sec	2880000000	2850000000	2890000000	15491933.3848297	0.00537914353639919	
tcpbench_S1000000_t10_10.3.0.33	PASS	PASS	PASS	PASS	PASS	unit <th>mean</th> <th>minimum</th> <th>maximum</th> <th>deviation</th> <th>relative</th> <th></th>	mean	minimum	maximum	deviation	relative	
sender	2689773625	2389946000	2698565750	2690709125	2691350750	bits/sec	2631069050	2389946000	2693565750	120567998.88778	0.0458247186206611	
tcpbench_S1000000_t10_n100_10.3.0.33	PASS	PASS	PASS	PASS	PASS	unit <th>mean</th> <th>minimum</th> <th>maximum</th> <th>deviation</th> <th>relative</th> <th></th>	mean	minimum	maximum	deviation	relative	
sender	2765055750	2450020625	2758966375	2758764750	2764942000	bits/sec	2699549900	2450020625	2765055750	124794806.118369	0.0462280049419975	
iperf3_c10.3.0.33_u_b10G_w1m_t10	PASS	PASS	PASS	PASS	PASS	unit <th>mean</th> <th>minimum</th> <th>maximum</th> <th>deviation</th> <th>relative</th> <th></th>	mean	minimum	maximum	deviation	relative	
sender	1180000000	1120000000	1170000000	1170000000	1180000000	bits/sec	1164000000	1120000000	1180000000	22449944.3206436	0.0192868937462574	
receiver	1160000000	1100000000	1150000000	1150000000	1150000000	bits/sec	1142000000	1100000000	1160000000	21354156.5040626	0.0186989111243981	
iperf3_c10.3.0.33_u_b10G_w1m_t10_R	PASS	PASS	PASS	PASS	PASS	unit <th>mean</th> <th>minimum</th> <th>maximum</th> <th>deviation</th> <th>relative</th> <th></th>	mean	minimum	maximum	deviation	relative	
sender	1230000000	1230000000	1230000000	1220000000	1230000000	bits/sec	1228000000	1220000000	1230000000	4000000	0.00325732899022801	
receiver	1230000000	1230000000	1230000000	1220000000	1240000000	bits/sec	1230000000	1220000000	1240000000	6324555.32033676	0.00514191489458273	
time_ip_make_CGENERIC.MP_j8_s	PASS	PASS	PASS	PASS	PASS	unit <th>mean</th> <th>minimum</th> <th>maximum</th> <th>deviation</th> <th>relative</th> <th></th>	mean	minimum	maximum	deviation	relative	
real	127.16	124.89	125.81	124.68	126.15	sec	125.738	124.68	127.16	0.898385218044017	0.0071448982649956	
user	433.74	432.79	432.72	433.14	434.53	sec	433.384	432.72	434.53	0.67724736987306	0.00156269583065609	
sys	392.39	395.47	394.06	393.57	394.87	sec	394.072	392.39	395.47	1.06512722244811	0.00270287465856014	
time_ip_fs_mark_dfs_mark_D8_N16_n256_t8	PASS	PASS	PASS	PASS	PASS	unit <th>mean</th> <th>minimum</th> <th>maximum</th> <th>deviation</th> <th>relative</th> <th></th>	mean	minimum	maximum	deviation	relative	
files	148.7	157.4	153.1	154.4	160.9	1/sec	154.9	148.7	160.9	4.10316950661316	0.0264891511078964	

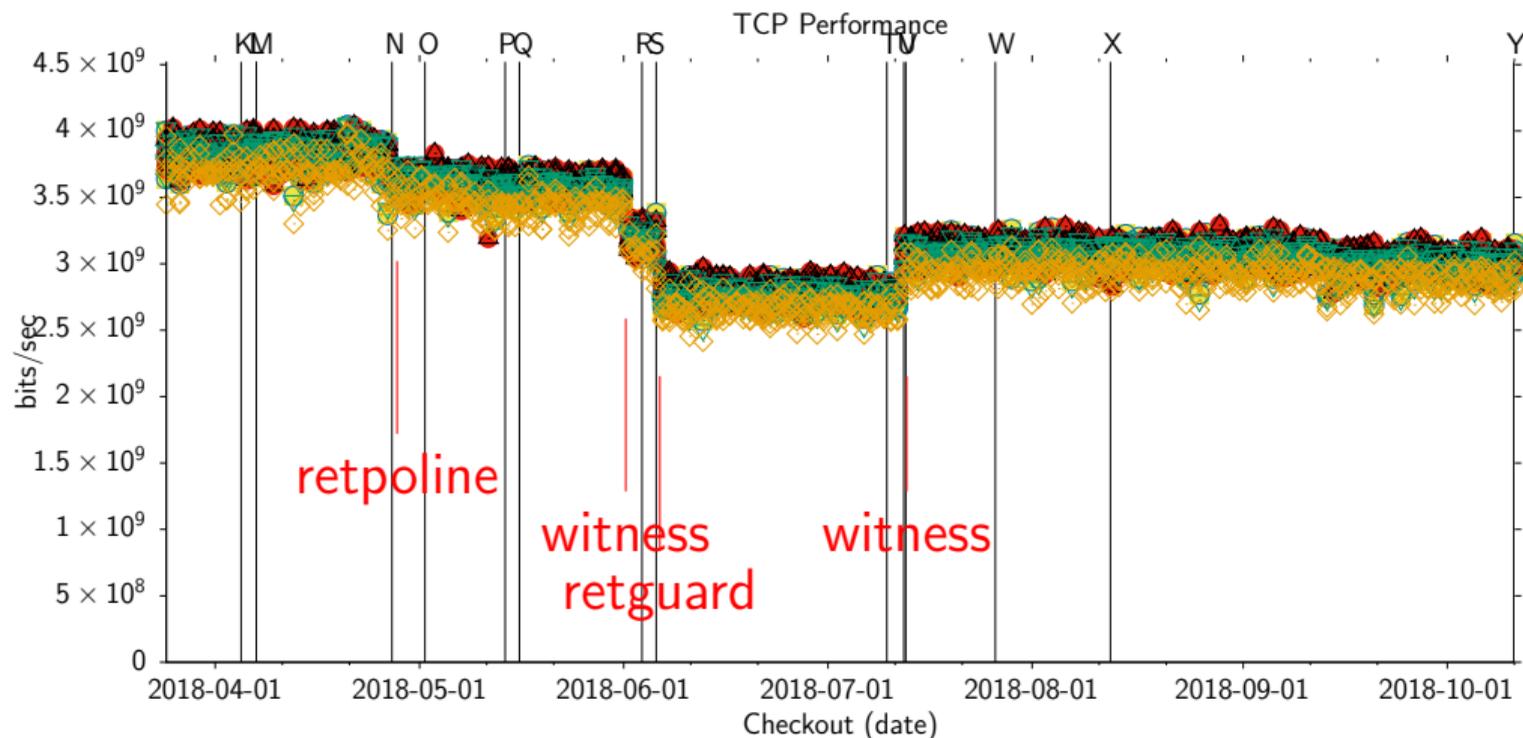
from 6.2 to 6.3, 173 Days, Reorder



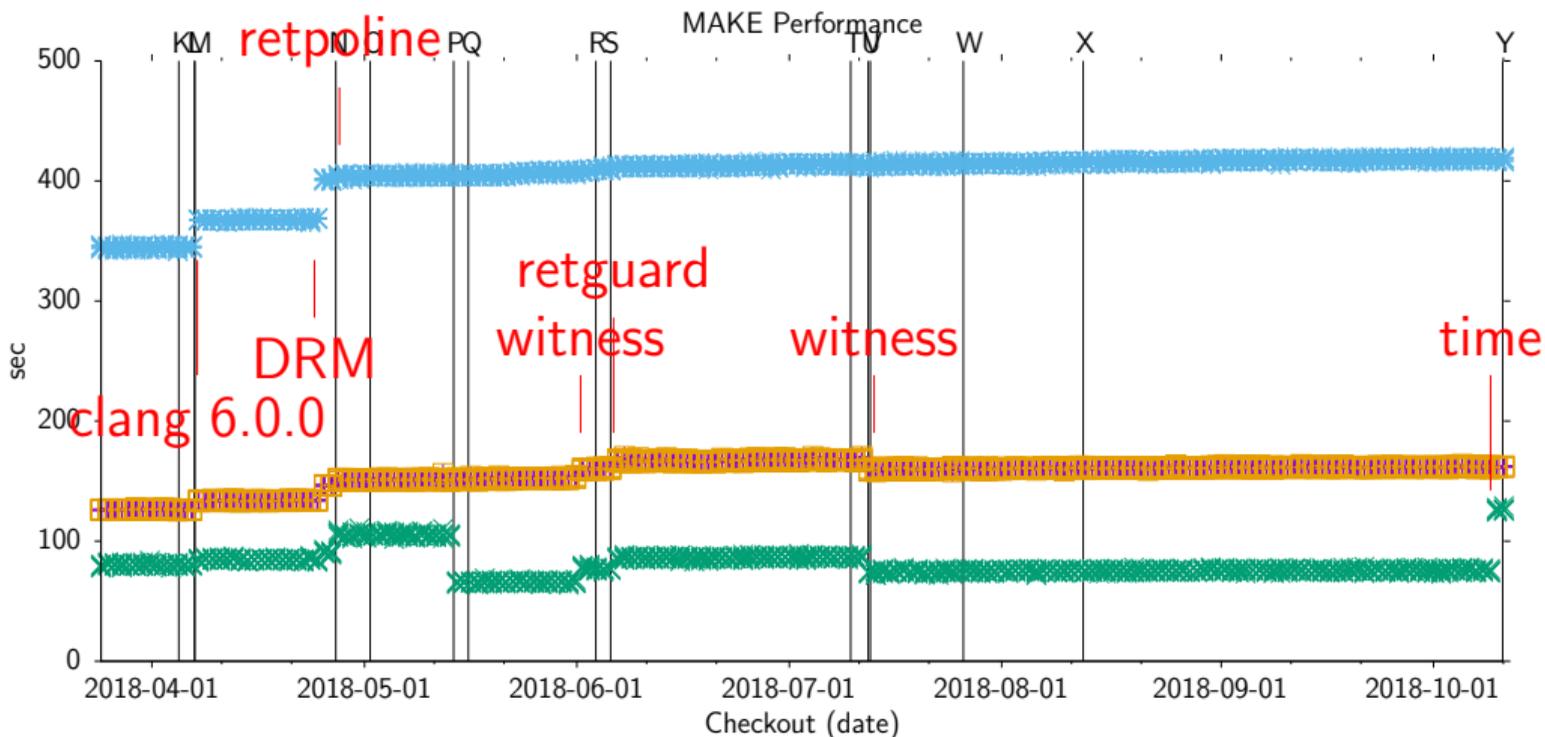
from 6.2 to 6.3, 173 Days, Make Kernel



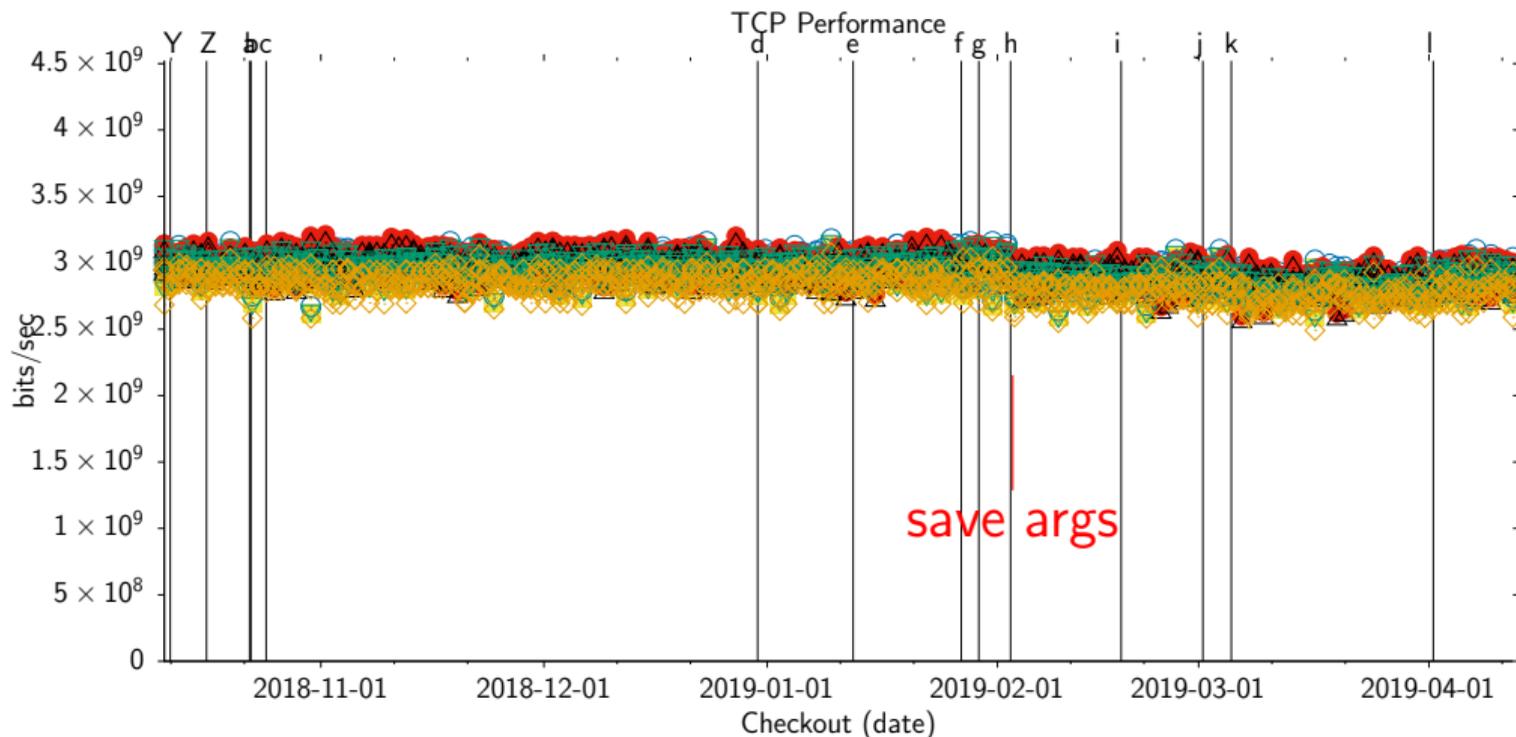
from 6.3 to 6.4, 202 Days, Reorder



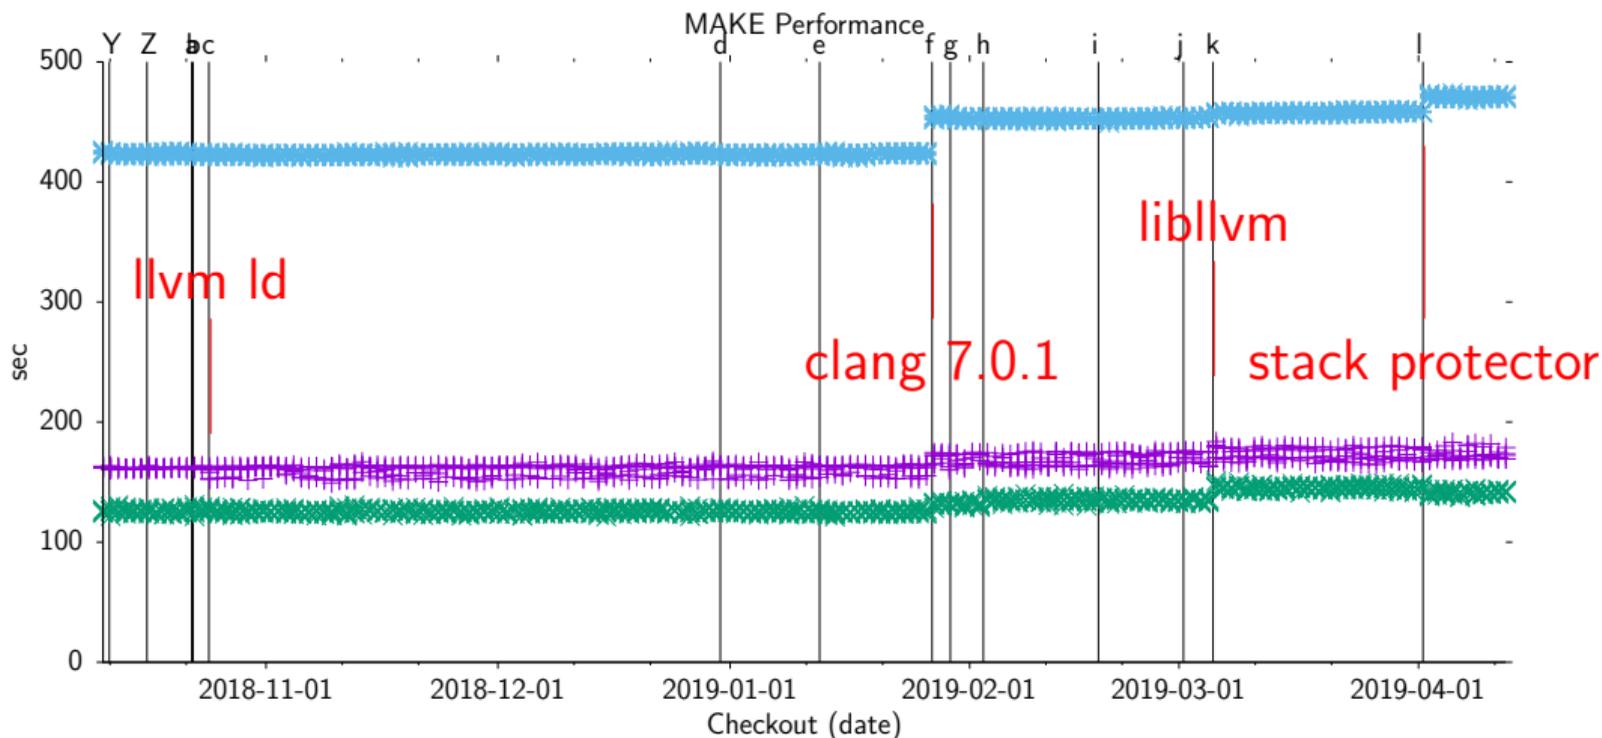
from 6.3 to 6.4, 202 Days, Reorder



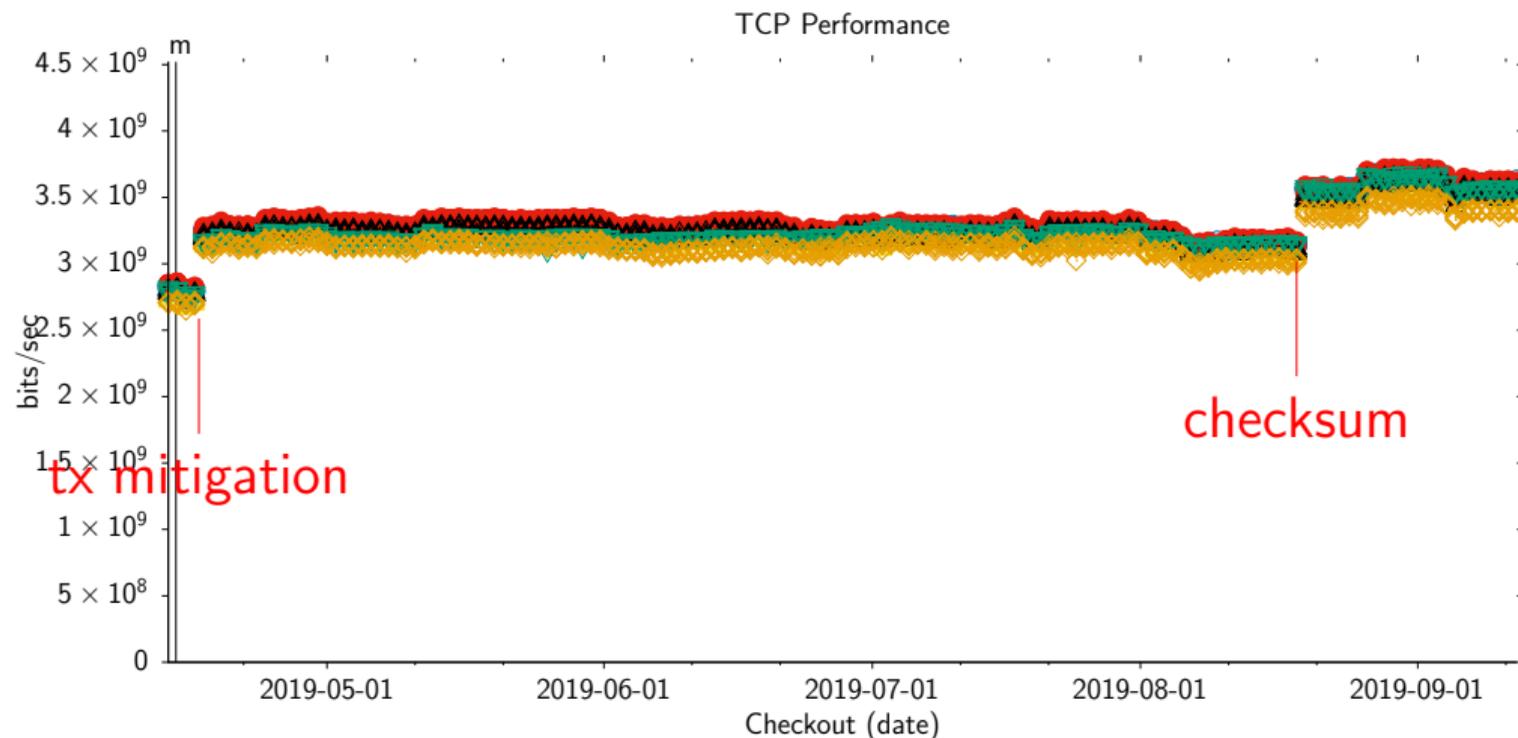
from 6.4 to 6.5, 185 Days, Reorder



from 6.4 to 6.5, 185 Days, Make Kernel



from 6.5, 154 Days, Align



OpenBSD CVS Log

created 2019-04-20T18:30:24Z

begin 2019-04-16T00:00:00Z

end 2019-04-17T00:00:00Z

path src/sys

commits 8

date 2019-04-16T04:04:19Z

author dlg

files src/sys/net/if.c [log](#) [diff](#) [annotate](#)

src/sys/net/if_var.h [log](#) [diff](#) [annotate](#)

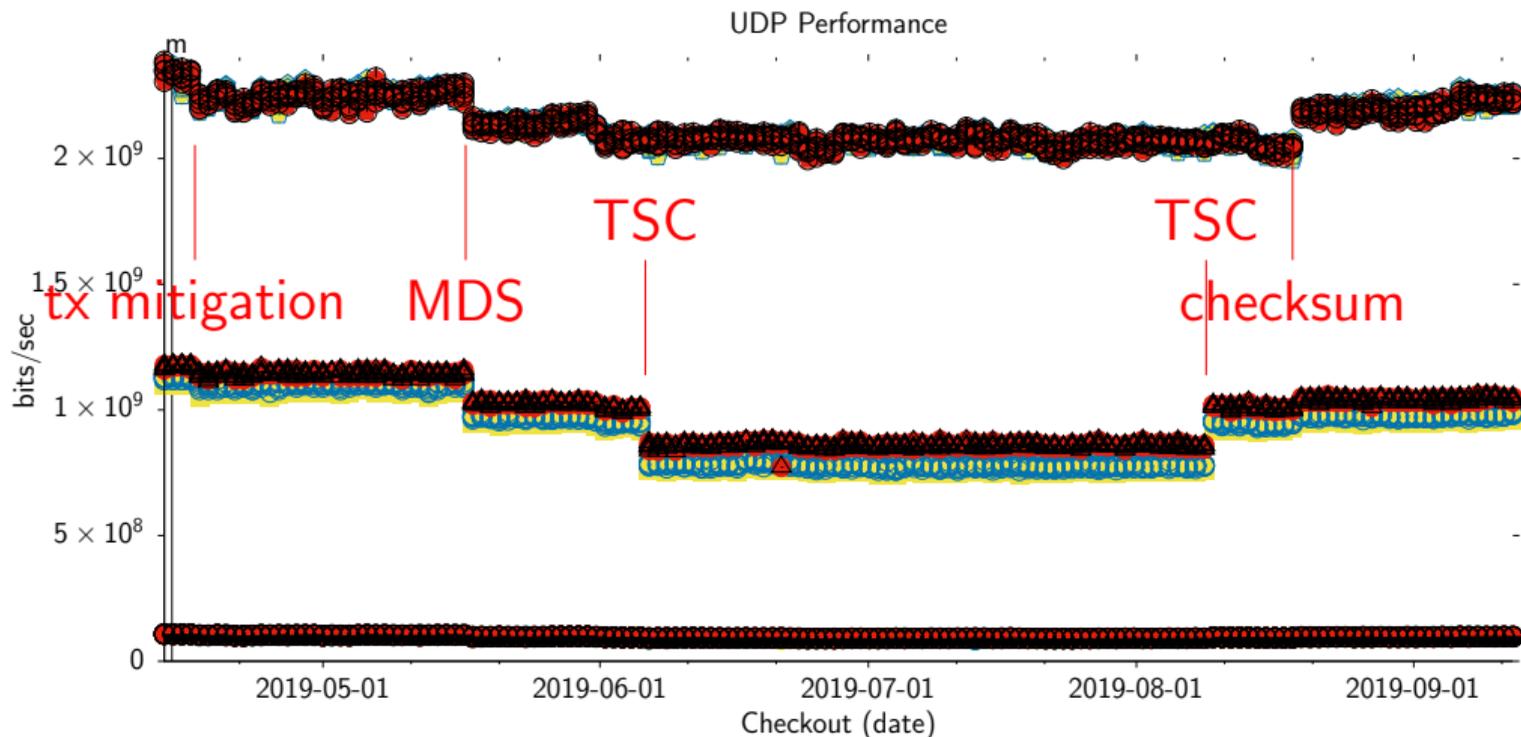
src/sys/net/ifq.c [log](#) [diff](#) [annotate](#)

src/sys/net/ifq.h [log](#) [diff](#) [annotate](#)

message have another go at tx mitigation

the idea is to call the hardware transmit routine less since in a lot of cases posting a producer ring update to the chip is (very) expensive. it's better to do it for several packets instead of each packet, hence calling this tx mitigation.

UDP Throughput, from 6.5, 154 Days, Align



UDP and Timecounter

iperf3

timecounter	UDP Mbits
tsc	924
acpihp0	739
acpitimer0	395
i8254	306

iperf3 UDP

send packet in iperf3 loop

- 1 write
- 2 gettimeofday
- 1 select
- 2 gettimeofday

Agenda

- 1 What did exist before?
- 2 How does it work?
- 3 What are the findings?
- 4 What is the Conclusion?**

Insights

- measuring sucks
- multi socket CPUs suck
- reproducing is hard
- do not trust your numbers
- keep it stupid simple

Future Ideas

- forwarding throughput
- Linux client and server
- testing patches
- historic releases
- file system performance

Thanks

- Jan Klemkow for Hardware Administration
- Moritz Buhl for Visualization
- genua for Hosting and Worktime

Links

- <http://bluhm.genua.de/>
- <http://bluhm.genua.de/regress/results/regress.html>
- <http://bluhm.genua.de/perform/results/perform.html>
- <http://bluhm.genua.de/perform/results/gnuplot/test.data>
- <https://github.com/bluhm/regress-all>
- <https://github.com/bluhm/udpbench>
- <https://github.com/younix/testmaster>
- <https://github.com/bluhm/talk-perform>

What did exist before?
○○○○○

How does it work?
○○○○○○○○○○

What are the findings?
○○○○○○○○○○○○○○○○○○○○○○○○○○○○○○

What is the Conclusion?
○○○○●

Questions

