

EuroBSDcon 2017 (FreeBSD devsumit)




# FreeBSD forwarding performance problems

Olivier Cochard-Labbé


# Queues/cores impact

Number of queues impact on forwarding performance  
(HP ProLiant DL360p Gen8 with 8 cores Intel Xeon E5-2650 2.60GHz, Chelsio T540-CR)



# Analysing bottleneck

```
kldload hwpmc  
pmcstat -S CPU_CLK_UNHALTED_CORE -l 20 -O data.out  
stackcollapse-pmc.pl data.out > data.stack  
flamegraph.pl data.stack > data.svg
```



lock on arpresolve

lock on fib4\_lookup

random\_harvest\_queue

NIC drivers  
& Ethernet path

# Already existing patches

- Random harvest problem

“Avoid spinning in random\_harvest\_queue”

<https://reviews.freebsd.org/D12132>

- fib4\_lookup (projects/routing)

“Convert RADIX and RIB lock to rmlock”

<https://people.freebsd.org/~ae/radix.diff>

- arp\_resolve (projects/routing)

“ Convert AFDATA lock to rmlock.”

<https://people.freebsd.org/~ae/afdata.diff>


# Yandex's patches

setup	11.1	11.1-Yandex	ministat
E5_2650-cxl Xeon & Chelsio NIC	5.75 Mpps	10.9 Mpps	90.56% +/- 1.24
E5_2650-oce Xeon & Emulex NIC	1.33 Mpps	1.33 Mpps	No diff. proven at 95.0% confidence
C2758-cxl Atom & Chelsio NIC	3.15 Mpps	4.2 Mpps	34.4% +/- 2.9%
C2758-ix Atom & Intel NIC	2.43 Mpps	3.08 Mpps	26% +/- 1.18
C2558-igb Atom & Intel NIC	1 Mpps	1.2 Mpps	20.17% +/- 2.56%
GX412-igb AMD & Intel NIC	747 Kpps	729 Kpps	-2.37% +/- 0.58%

harvest\_mask="351"

# Linear performance ?

Number of queues impact on forwarding performance  
(HP ProLiant DL360p Gen8 with 8 cores Intel Xeon E5-2650 2.60GHz, Chelsio T540-CR)



# inet4 vs inet6 performance

setup	inet4 (median)	inet6 (median)	ministat
E5_2650-cxl Xeon & Chelsio NIC	10.94 Mpps	9.18 Mpps	-16.12% +/- 0.19%
C2758-cxl Atom & Chelsio NIC	4.29 Mpps	3.43 Mpps	-19.08% +/- 1.61%
C2758-ix Atom & Intel NIC	3.81 Mpps	3.43 Mpps	-9.84% +/- 1.3%
C2558-igb Atom & Intel NIC	1.23 Mpps	1.08 Mpps	-11.79% +/- 0.5%
GX412-igb AMD & Intel NIC	734 Kpps	709 Kpps	-3.6% +/- 0.70%

# VLAN tagging

x Xeon E5\_2650-cx1, no VLAN tagging: inet4 packets-per-second

+ Xeon E5\_2650-cx1, VLAN tagging: inet4 packets-per-second

\* Xeon E5\_2650-cx1, VLAN tagging and no hardware acceleration: inet4 pps



	N	Min	Max	Median	Avg	Stddev
x	5	10917371	10970686	10945136	10946743	22298.313
+	5	9056449	9104195	9064032	9075563.7	21531.387

Difference at 95.0% confidence  
-1.87118e+06 +/- 31966.4  
-17.0935% +/- 0.267353%  
(Student's t, pooled s = 21918.2)

*	5	8007220	8074732	8017098	8029088.6	26769.461
---	---	---------	---------	---------	-----------	-----------

Difference at 95.0% confidence  
-2.91765e+06 +/- 35929.5  
-26.6532% +/- 0.295534%  
(Student's t, pooled s = 24635.5)

-17% with tagging: Known problem  
Yet another patch from Yandex  
<https://reviews.freebsd.org/D12040>





# All online

- Configuration files, ministat result and flamegraph
- <https://github.com/occochard/netbenches>
-