

Classifier Project

Olof Hagsand, Bengt Görden,
Jens Låås, Robert Olsson

Funded by IIS

Classifier Project

Work-In-Progress

Machine TYAN S7025 just arrived
Intel 82599 NIC's in December from Intel

Classifier Project

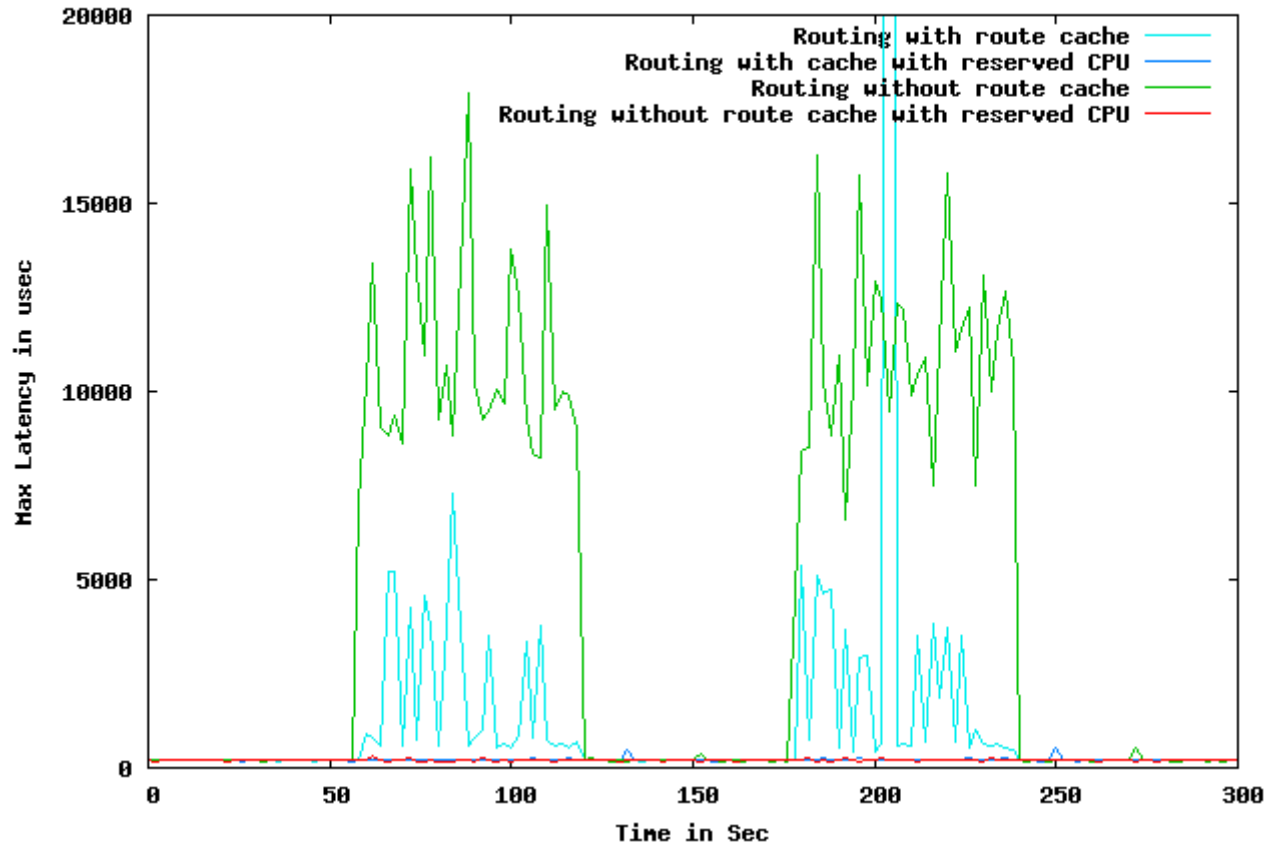
General Question:

Can we improve robustness
and performance with the new
NIC HW-based classifiers?

How?

Effects of packet load

Latency comparison between hard and relatively moderate routing load
Userland Latency Response On CPU0 when routing with and without route cache
load routed under 60-120 resp. 180-240 sec in both cases



User space scheduling latency under heavy network load

What can be done?

We have many queues/CPU's.
Can we reserve for ssh and bgp etc?

With interrupt affinity we can leave one
CPU free to run the daemons.

A “reservation” experiment

Measuring TCP transaction performance
on an idle interface.


During different irq affinity setups.

Reserve by CPU by affinity.

Using taskset for daemon(netserver).


Will be described here... :)

How much can the box sustain?
10 10g ports or 5x2



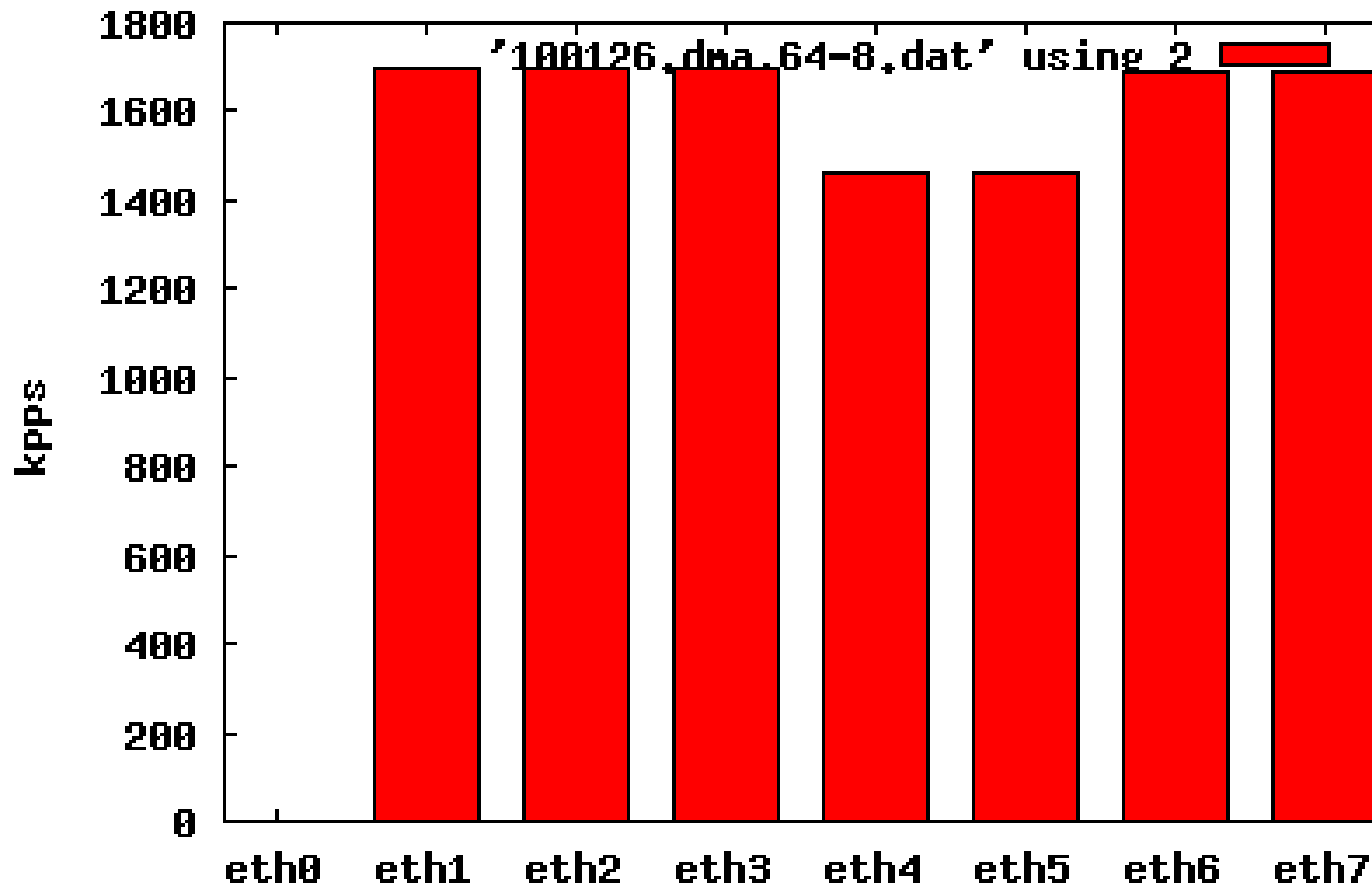
First shot approx 70 Gbit/s we need more CPU's :)

1500 bytes packets TX performance
Not to bad, We approach 70 Gbit/s...
And eth8 and eth9 is unused



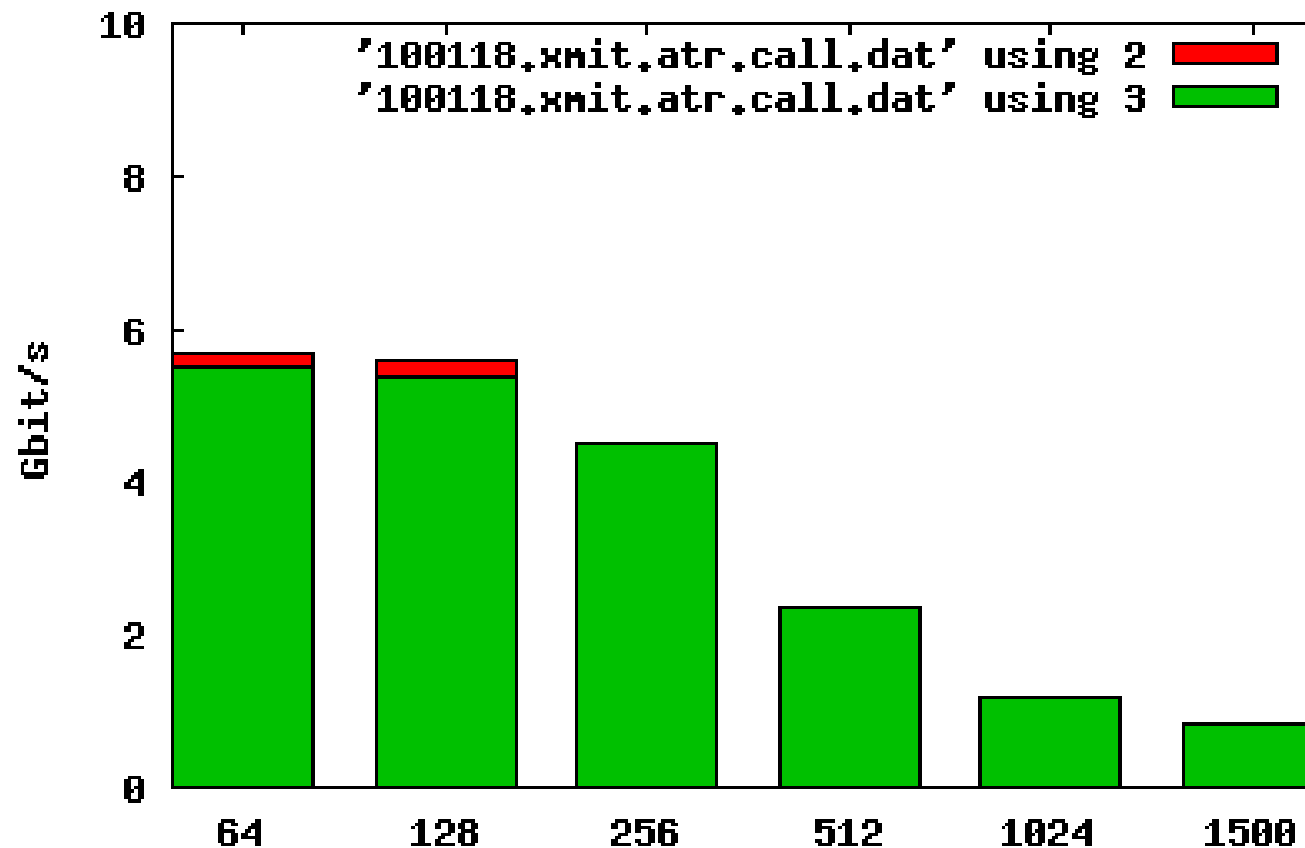
Weird? Small pkts

64 bytes packets TX performance
Things needs to be investigated...
CPU limitation?



TX cost of Flow Director

Flow director TX cost
New things for Intel 82599 chips...
Not to bad



Last a very promising experiment

- 1) We hacked the ixgbe driver to direct all ssh and bgp to CPU0
- 2) We managed remove the bulk data (RSS) from CPU0.

As result the “control-plane” now separated

Disclaimer needs to be further studied!

More to come

Questions?