



Hardware Rate Limiting Control

Jesse Brandeburg

John Fastabend

Intel® Ethernet

Problem

- Intel hardware has hardware rate limiters
 - 82599 has per queue rate limit
 - Expressed as a divider of the link speed
 - Guarantees maximum rate of the queue
- Linux has no direct control or interface for controlling these rate limiters

Proposed Solution

- John Fastabend designed a sysfs interface to control the tx-queue output speed
- Directory exists today (via RPS code), adds a new file
 - Example from today:
 - /sys/class/net/p1p1/queues/rx-0/rps_cpus
 - Proposed
 - /sys/class/net/p1p1/queues/tx-0/rate_limit

Code

- We have some, you can download it....
- <https://github.com/jrfastab/tx-rate-limits>

Summary

- Problem

- Lots of networking hardware has transmit rate limiters
- Linux kernel doesn't have interface to the above

- Proposal

- Use sysfs to control each queue's transmit max rate

- Discussion

- Is this solution worth submitting to the kernel?
- What else **could** we do?
- Can you help?

