10 Gbit Hardware Packet Filtering Using Commodity Network Adapters

Luca Deri <deri@ntop.org>
Joseph Gasparakis <joseph.gasparakis@intel.com>

RIPE 61 - 15/11/2010

10 Gbit Monitoring Challenges [1/2]

- High number of packets to be analyzed (10 times as much as 1 Gbit).
- CPU-based traffic analysis is not feasible at these speeds as it will result in severe packet loss.
- Packet filtering is very important, in particular on WANs, in order to early discard those packets that are not supposed to be analyzed.

10 Gbit Monitoring Challenges [2/2]

- Operating systems handle 10 Gbit adapters as legacy 10 Mbit adapters (use ethX for any speed).

- Modern computing architectures are grounded on multicore, where multiple threads of execution process data concurrently.

- The outcome is that basically only one core can handle incoming traffic.

OS Networking Limitations

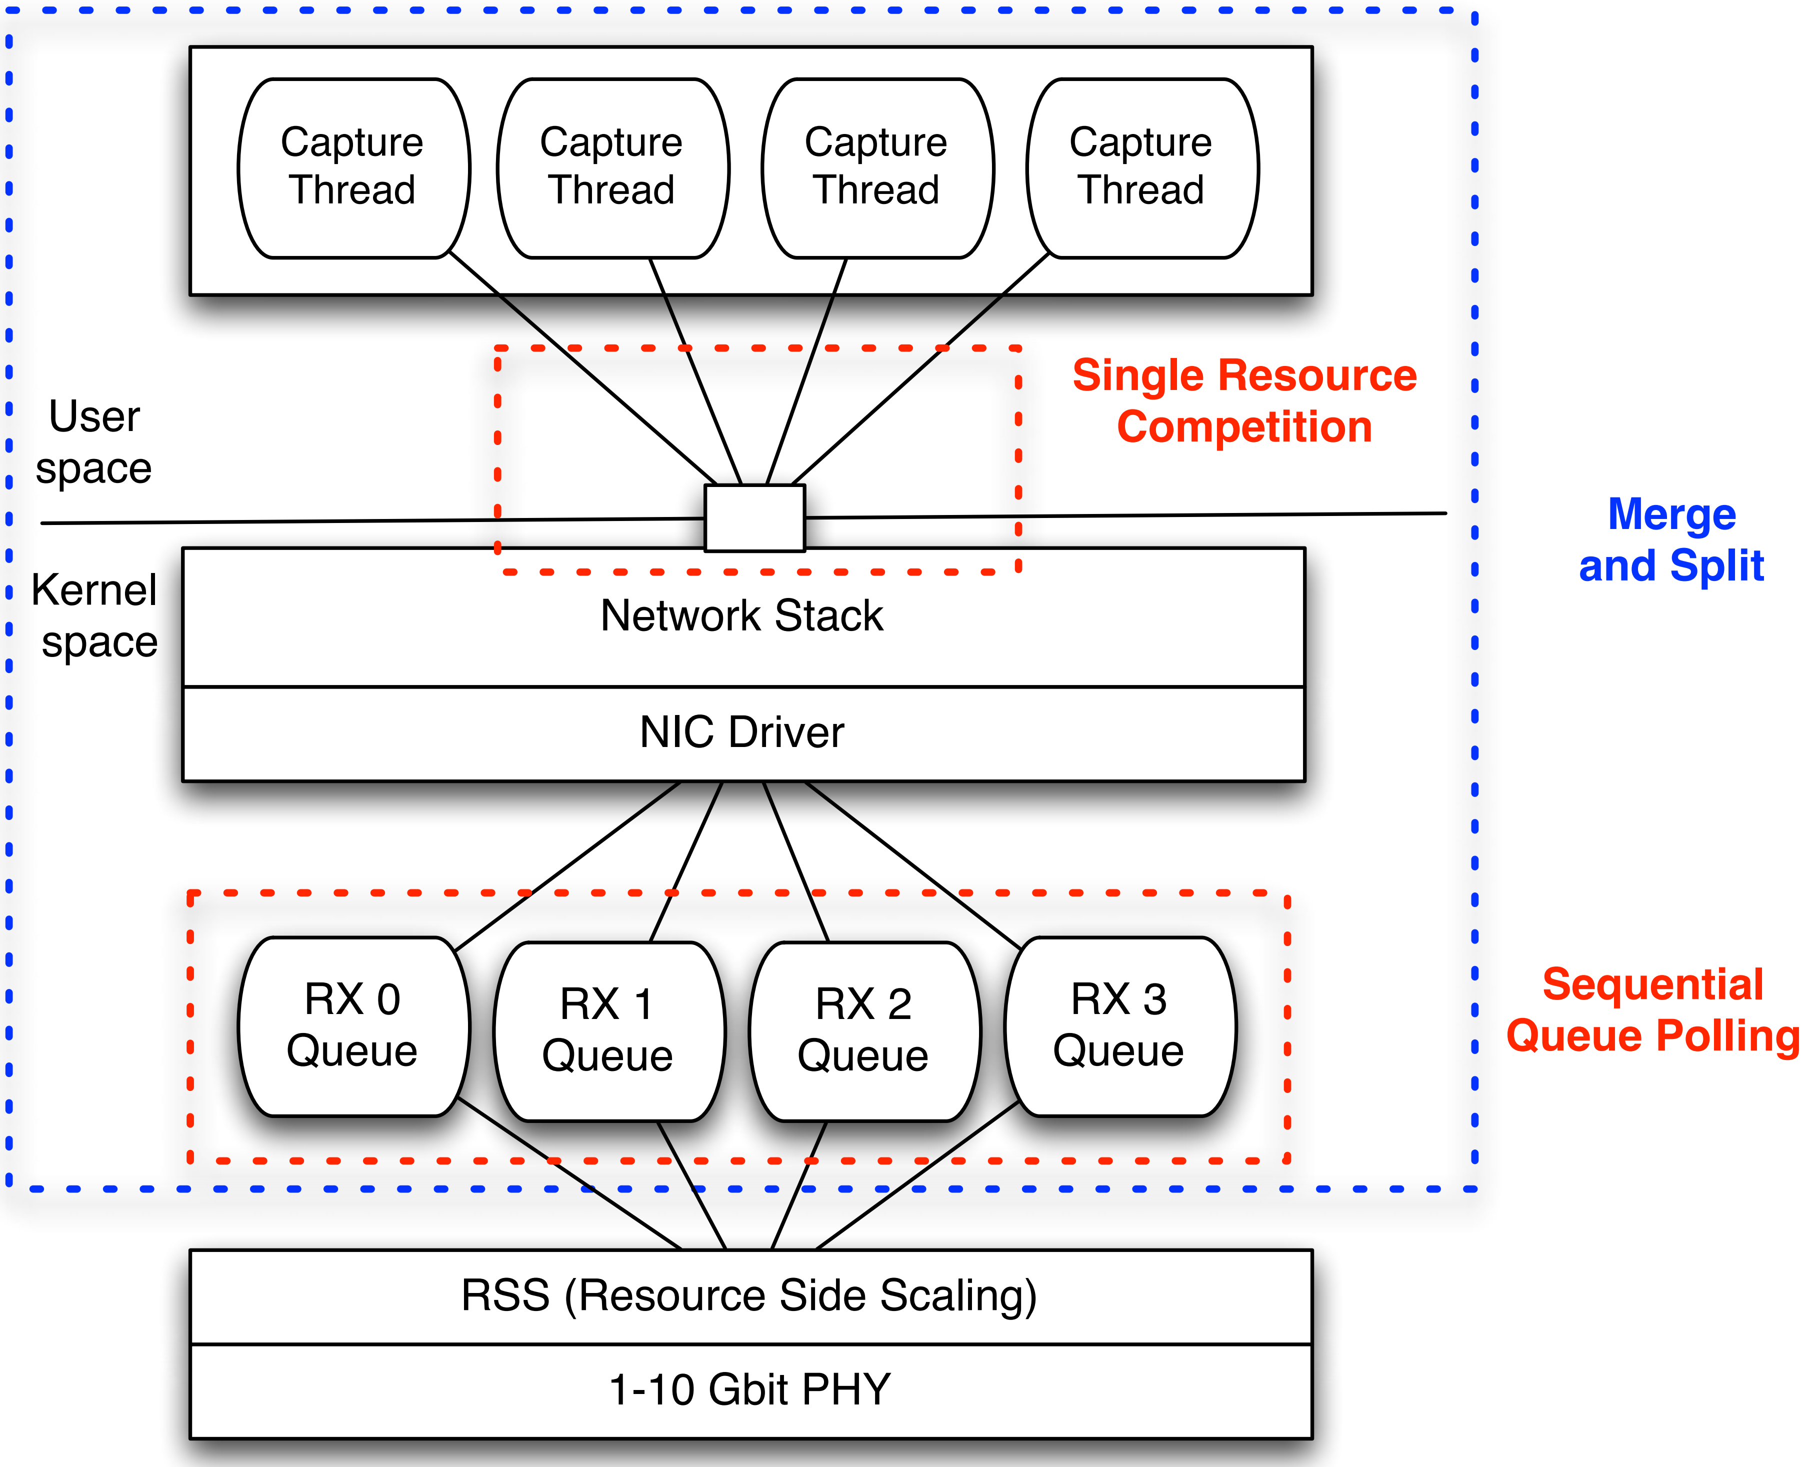

Hardware-based NICs

- FPGA-based Network Adapters
  - Endace DAG
  - Napatech
- Pros: ability to operate at wire rate
- Cons
  - High cost (> 10k USD)
  - Limited number of filtering rules (~32)

What about OpenFlow ?

- Network protocol that allows to remotely control the forwarding plane in switches

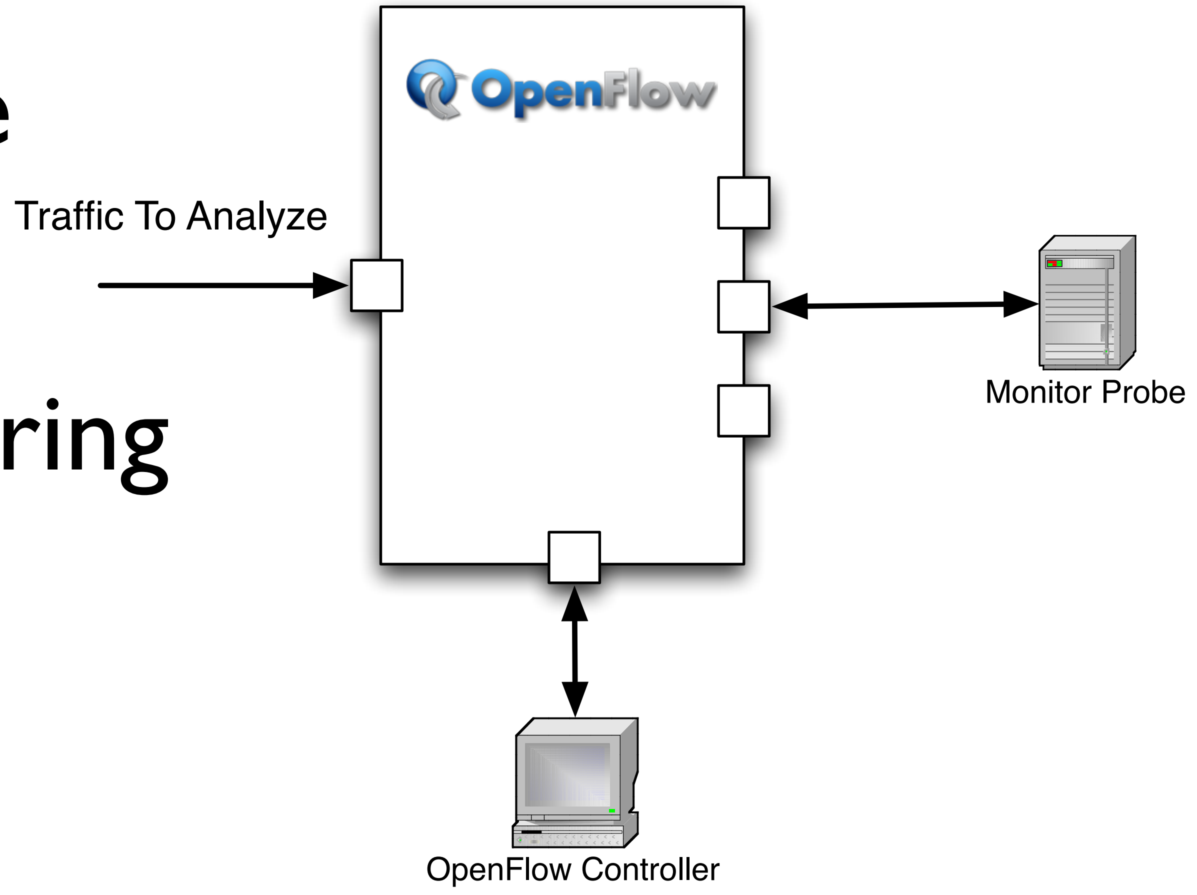

- Pros: Moves filtering in switches
- Cons:
  - 10 Gbit OpenFlow switches are costly
  - Complex architecture (cables & wires)

Why is Hw Filtering Important?

- It prevents unwanted traffic to reach the computer hence to waste CPU cycles.
- Filtering in software can lead to packet loss, thus having a negative drawback on analysis.
- Packet filtering is the cornerstone of efficiently dispatching incoming packets to available cores, that it's the only way to exploit modern computing architectures.

Modern Networking Architectures

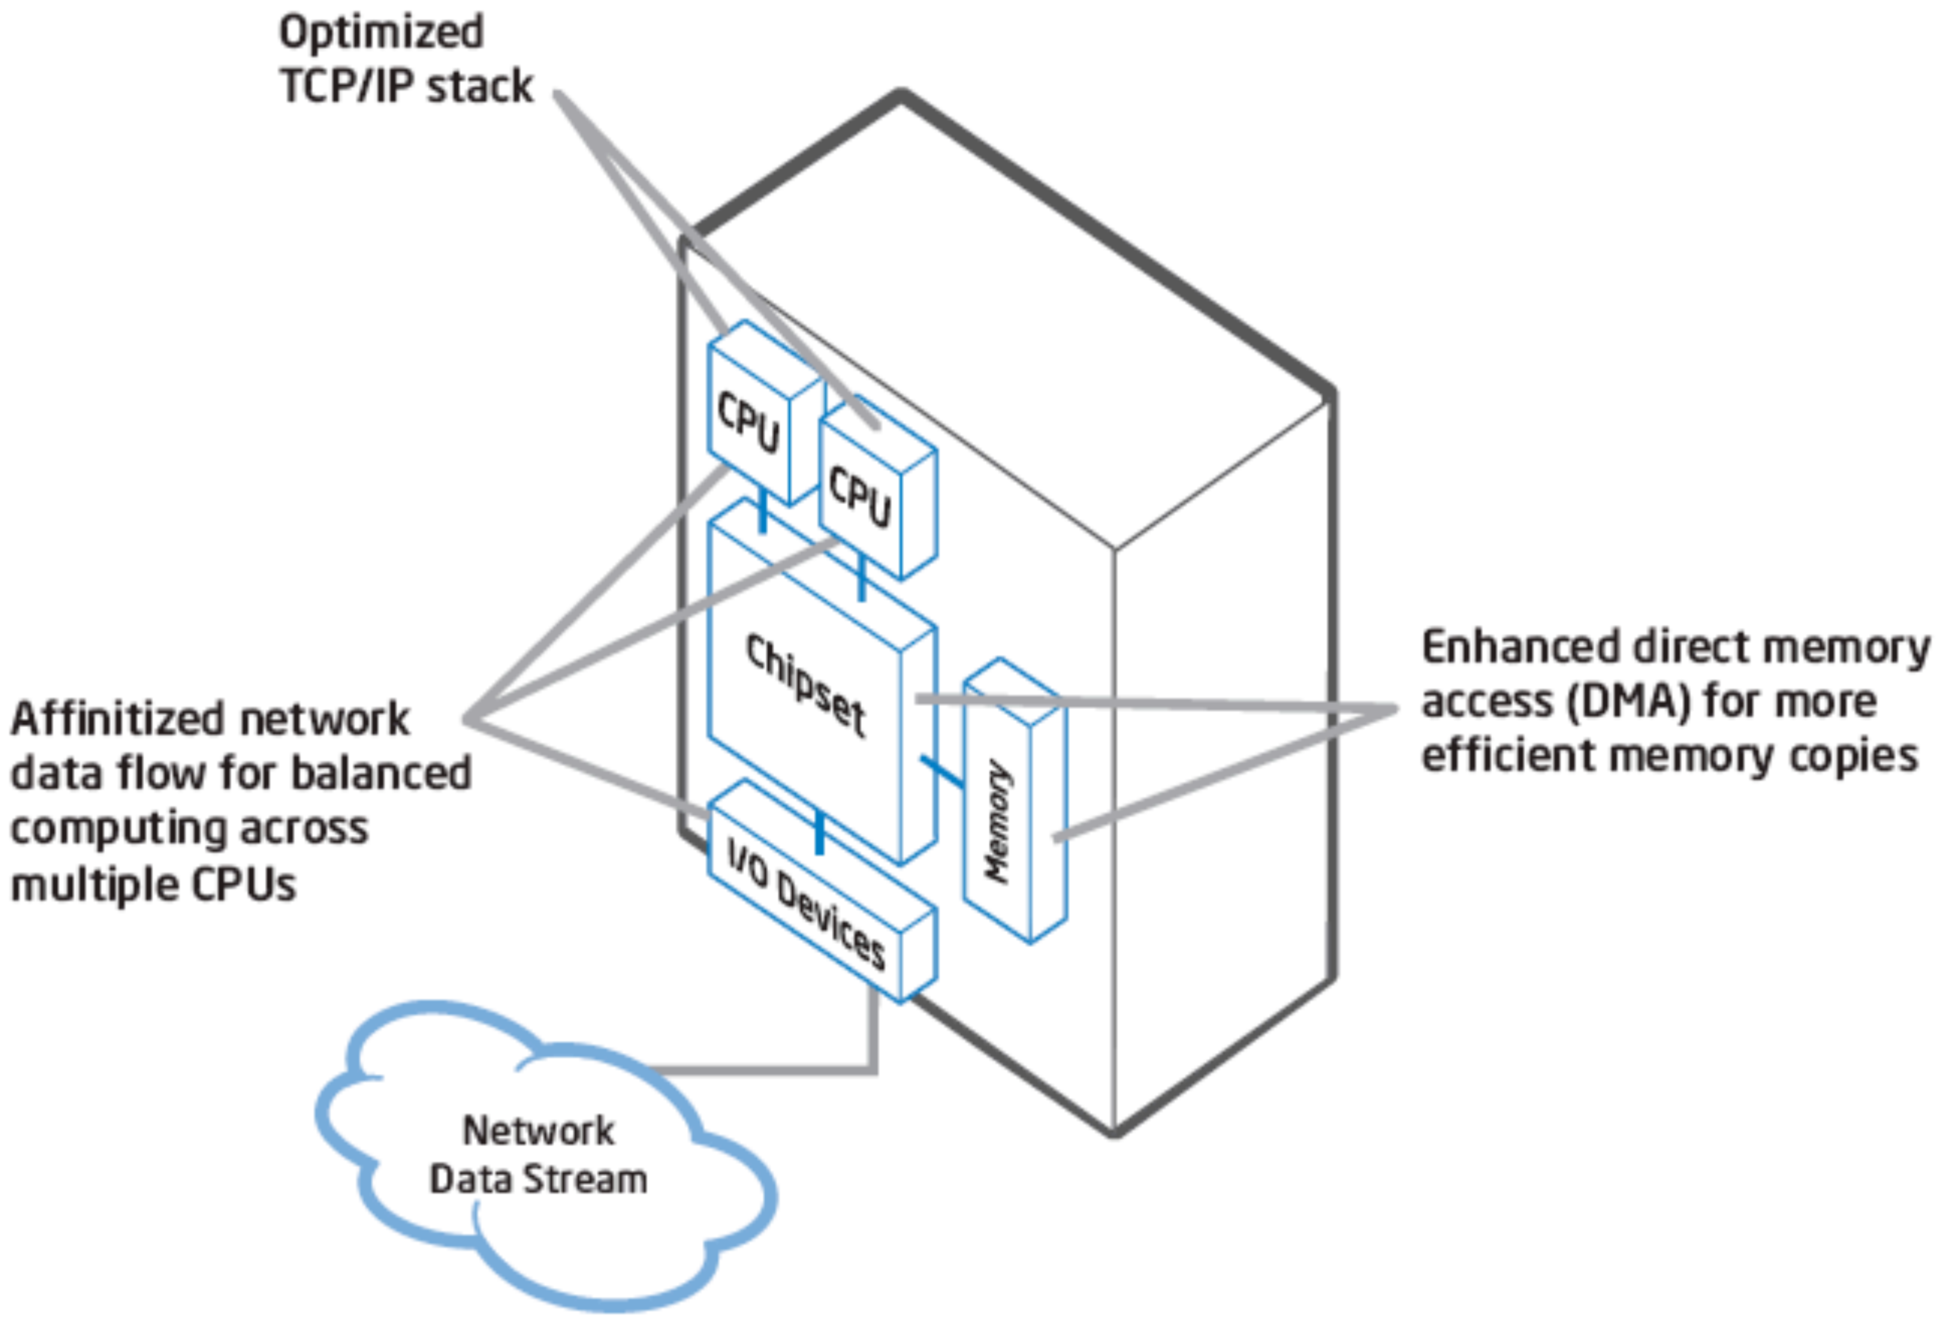

Optimized TCP/IP stack

Affinitized network data flow for balanced computing across multiple CPUs

Enhanced direct memory access (DMA) for more efficient memory copies

PF_RING+TNAPI

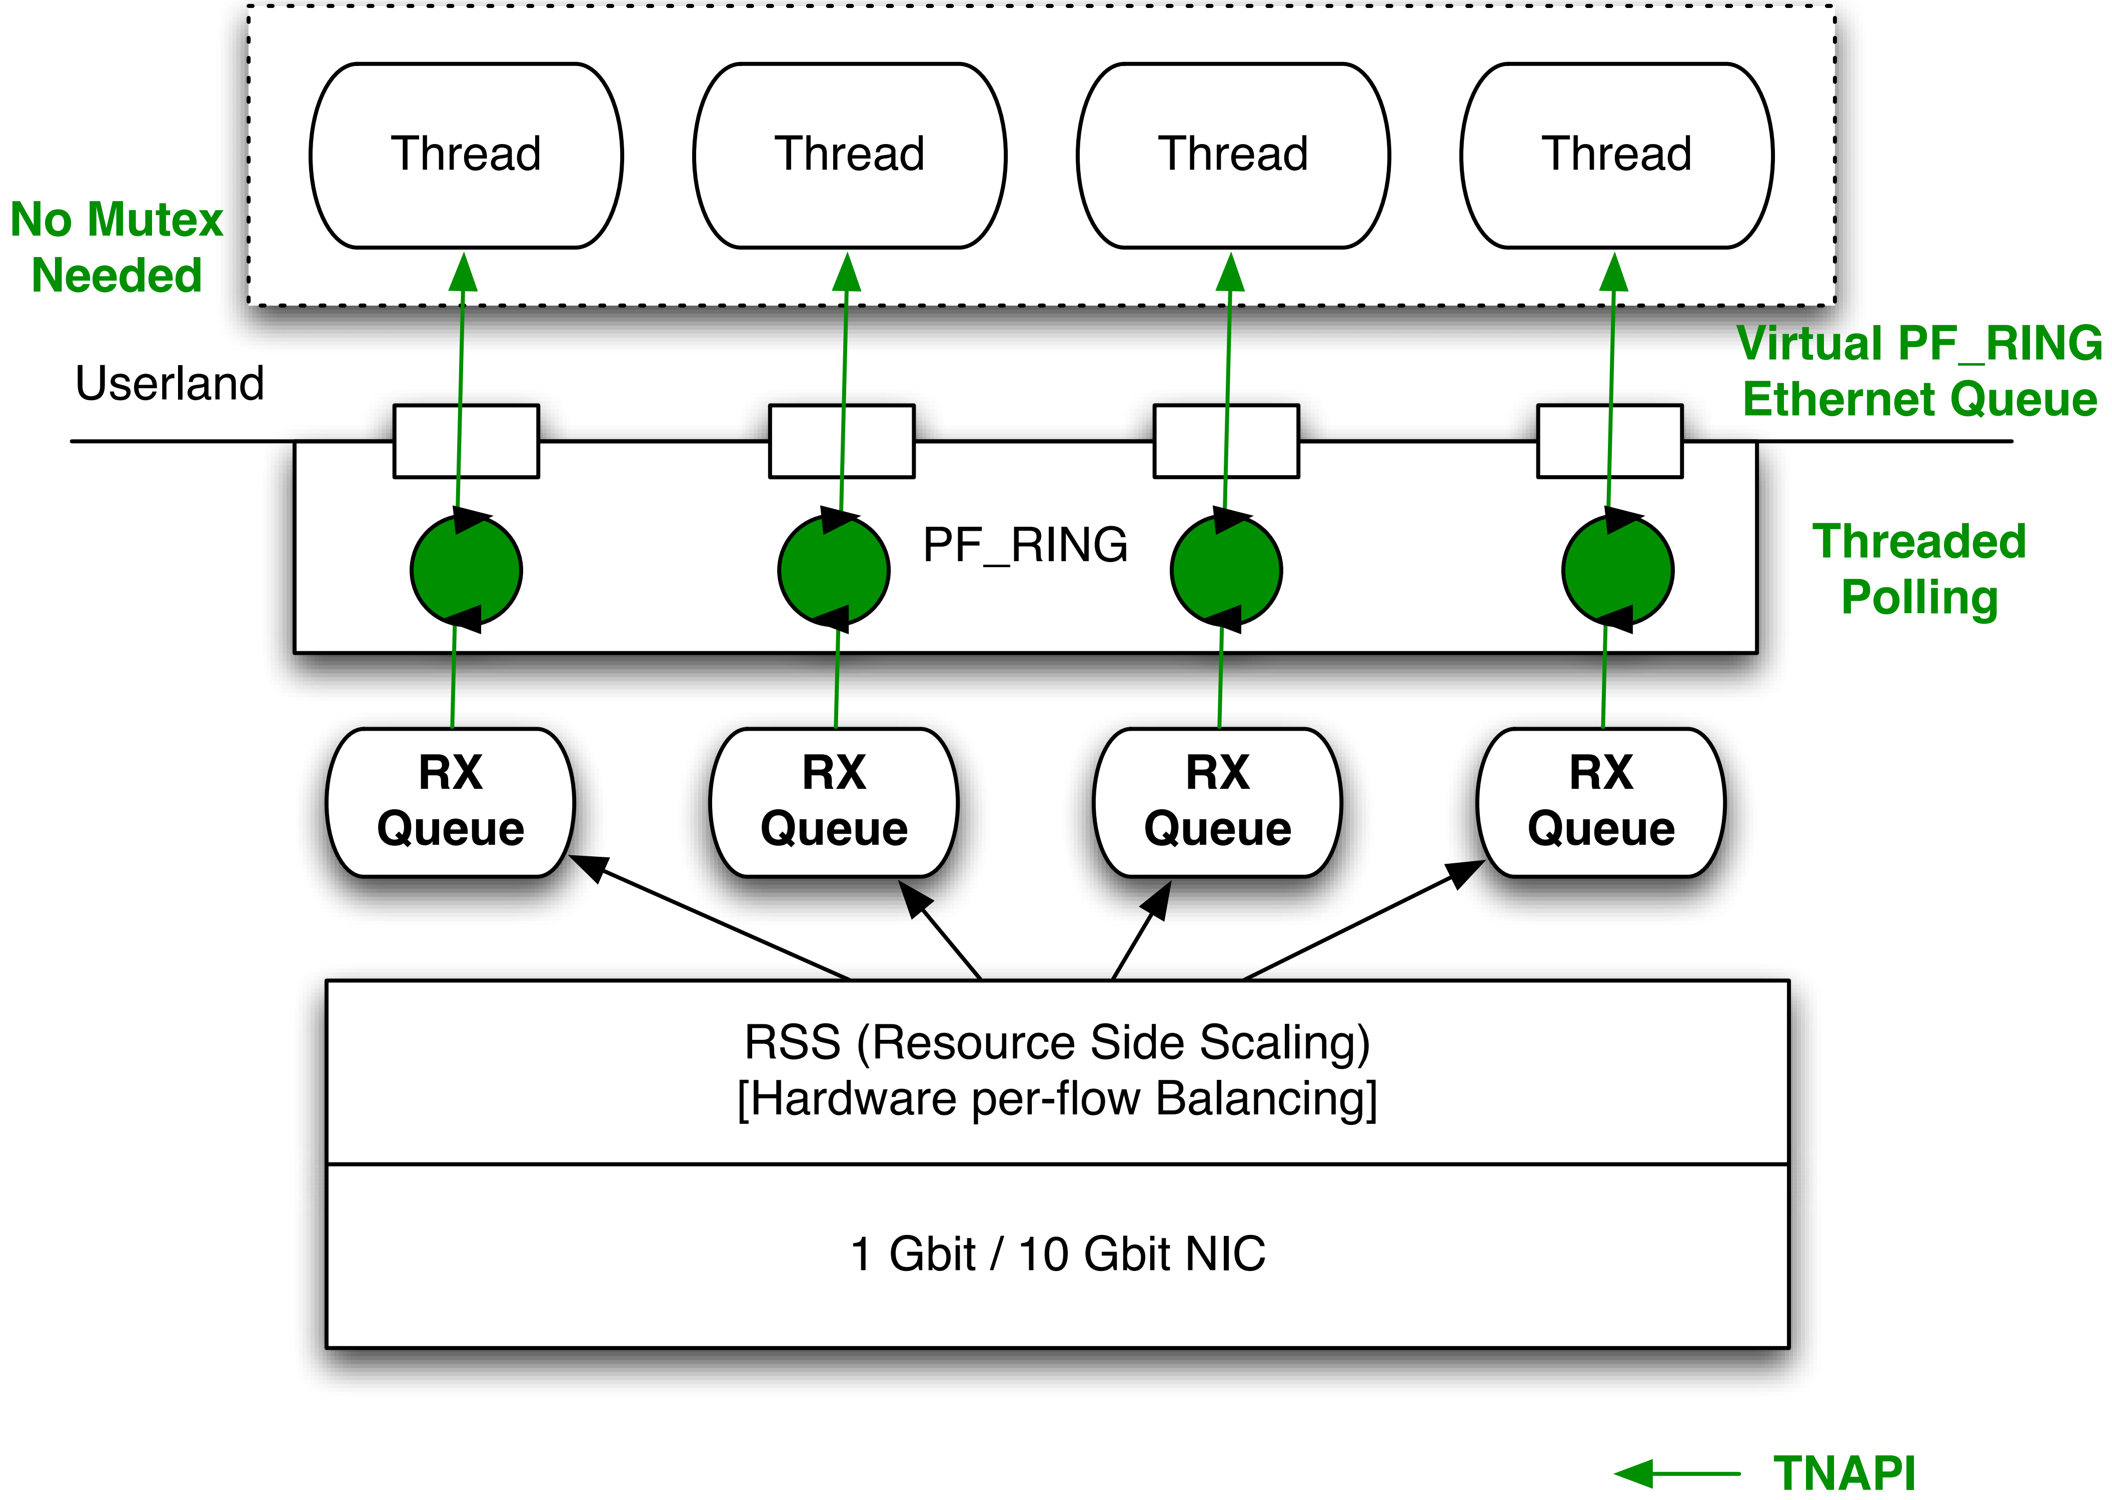

Intel 82599 Ethernet Controller [1/3]

- Latest generation of Intel 10 Gbit Ethernet Controller.
- Ability do define up to 32’000 perfect rules per port (unlimited hashing rules).
- Commodity adapter (<350 USD/port).
- Hardware support for virtualization (i.e. in-NIC L2 Switch) and multi RX/TX queues.
- Limitation: OSs exploits only basic NIC capabilities.

Intel 82599 Ethernet Controller [2/3]

- In 82599 packet filtering is performed in hardware at wire rate.

- Filtering is necessary to decide to which RX queue a packet must be assigned.

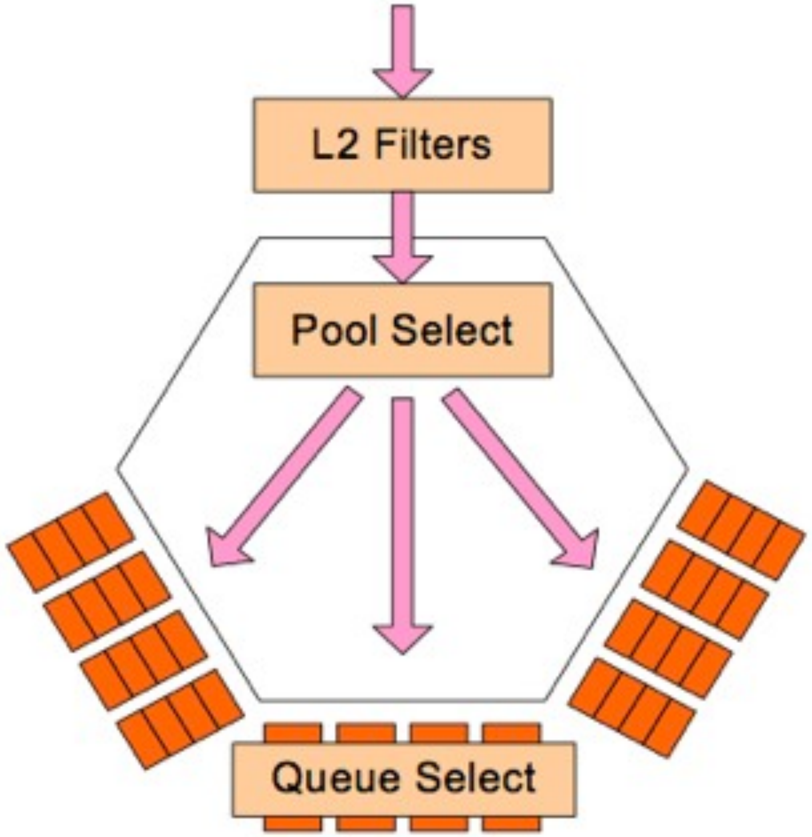

- Assigning a packet to a non-existing RX queue (<= number of available CPU cores) drops the packet.

Intel 82599 Ethernet Controller [3/3]

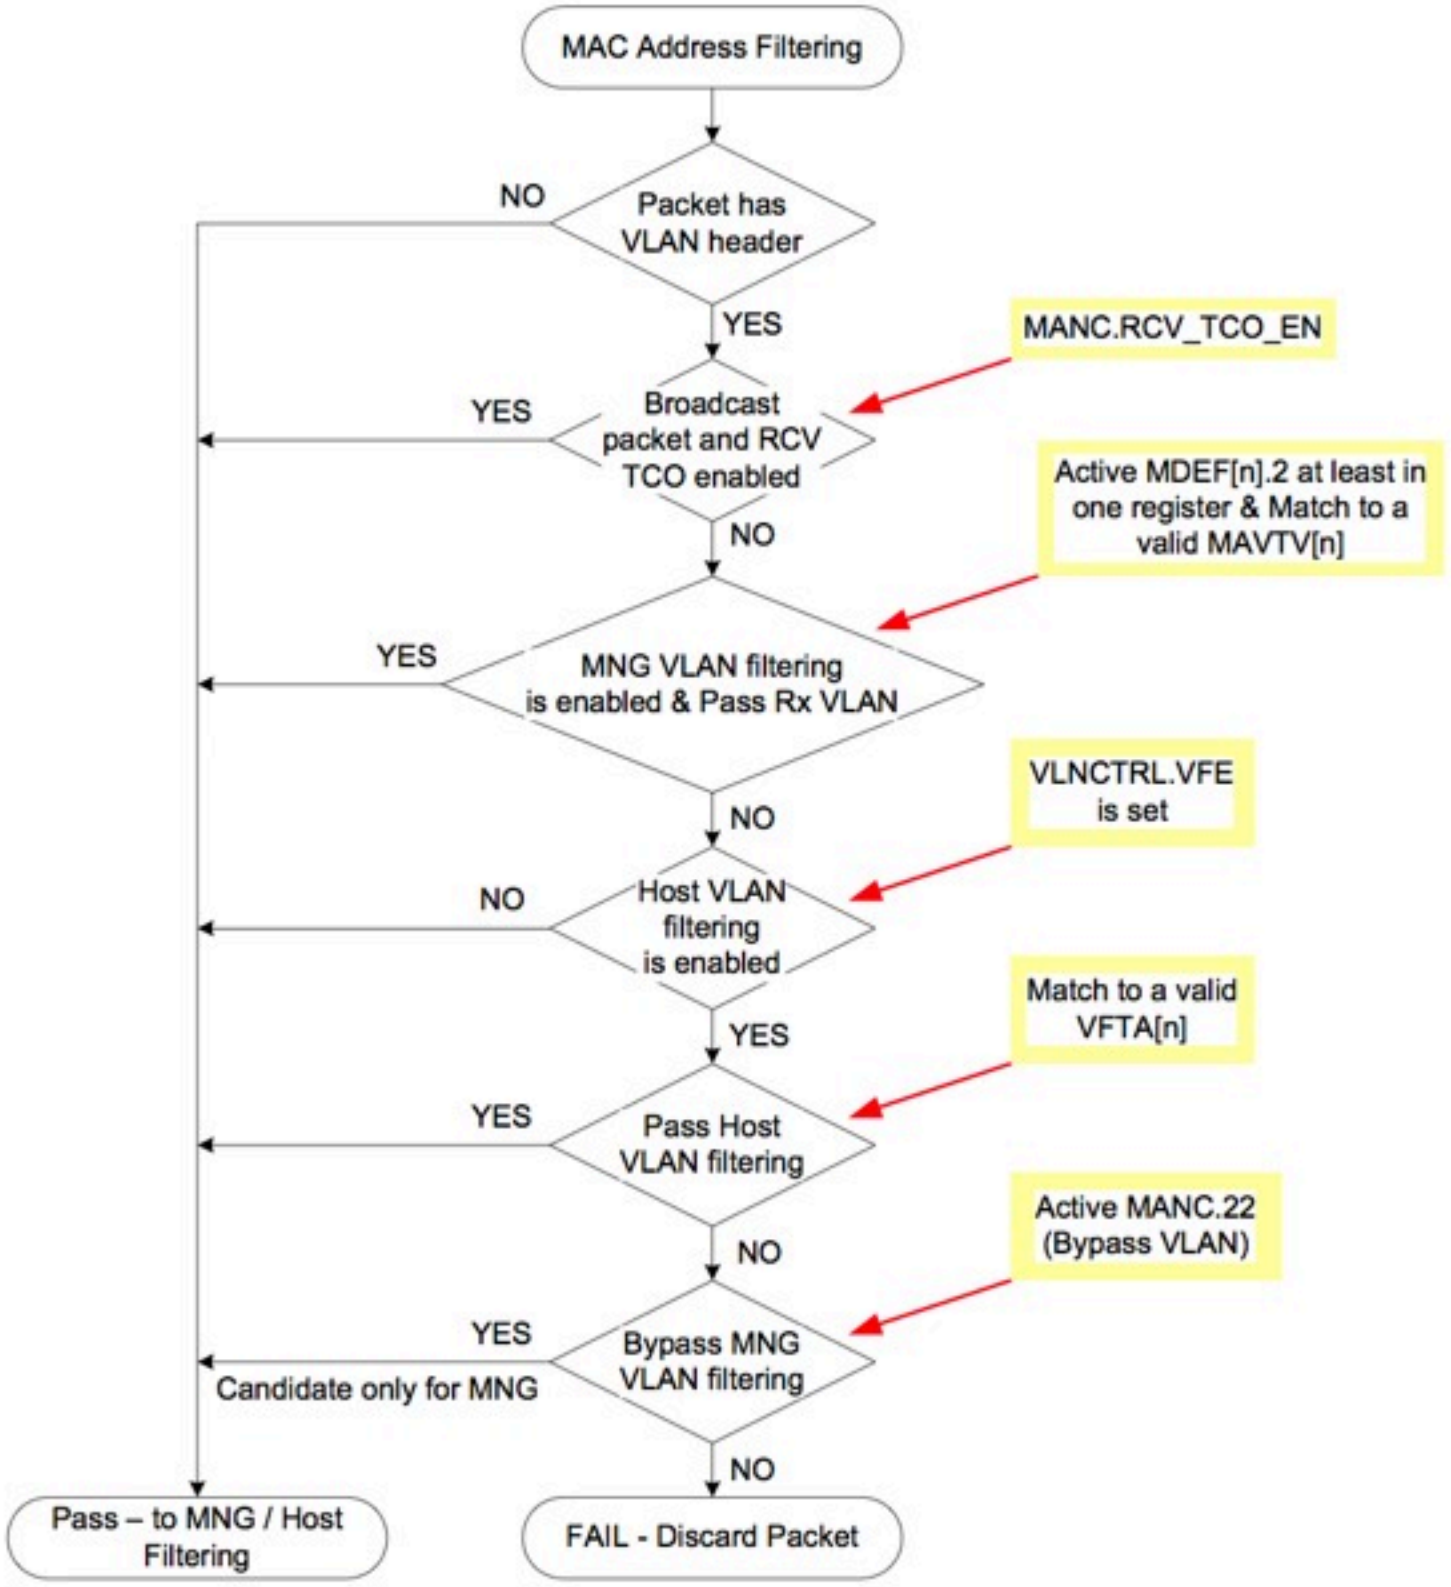

MAC Address Filtering
Packet has VLAN header
NO
YES
Broadcast packet and RCV TCO enabled
YES
NO
MNG VLAN filtering is enabled & Pass Rx VLAN
YES
NO
Host VLAN filtering is enabled
NO
YES
Pass Host VLAN filtering
YES
NO
Bypass MNG VLAN filtering
YES
Candidate only for MNG
NO
Pass – to MNG / Host Filtering
FAIL - Discard Packet
MANC.RCV_TCO_EN
Active MDEF[n].2 at least in one register & Match to a valid MAVTV[n]
VLNCTRL.VFE is set
Match to a valid VFTA[n]
Active MANC.22 (Bypass VLAN)

PF_RING with 82599 Support

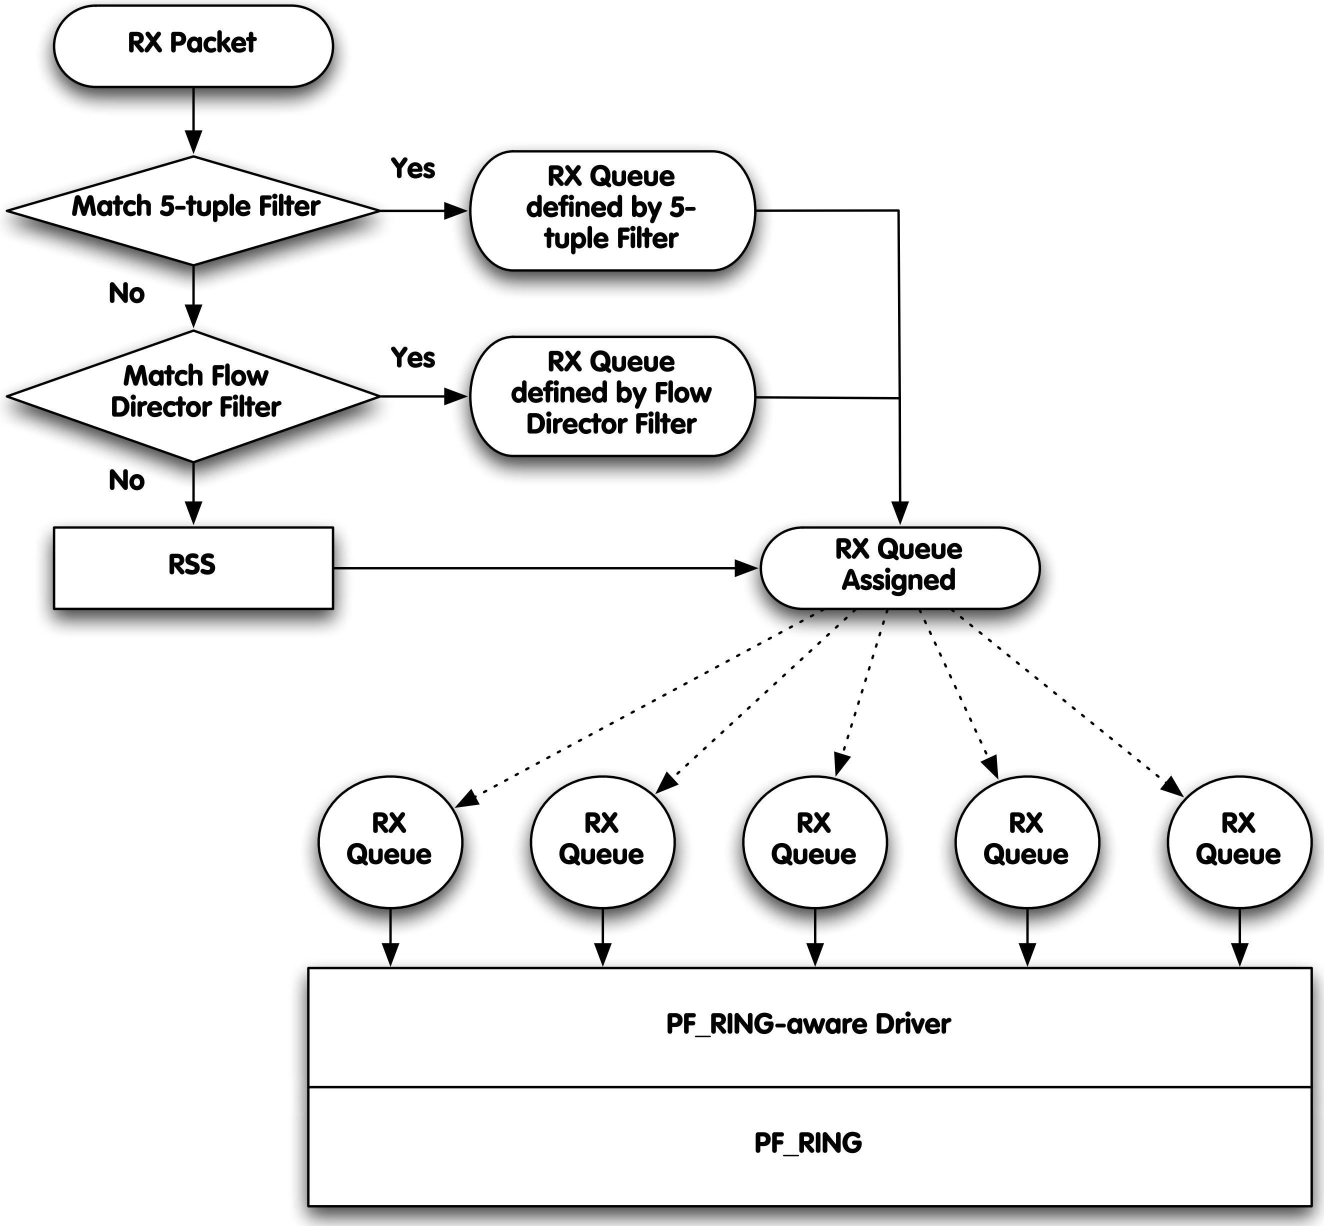

- `# echo "+(1,1,1,tcp,192.168.0.10/32,25,10.6.0.0/16,0)" > proc/net/pf_ring/dev/eth2/rules`
- `# echo "+(2,2,tcp,0.0.0.0,25,10.6.0.0,0)" > proc/net/pf_ring/dev/eth2/rules`

Using 82559 Filters in Real Life

- Signaling-based realtime multimedia (e.g. VoIP, IPTV) monitoring.
- Network Troubleshooting: Wireshark.
- Traffic Classification and Balancing.
- Lawful Interception of IP Traffic.
- 10 Gbit Firewalling.

10 Gbit Snorting

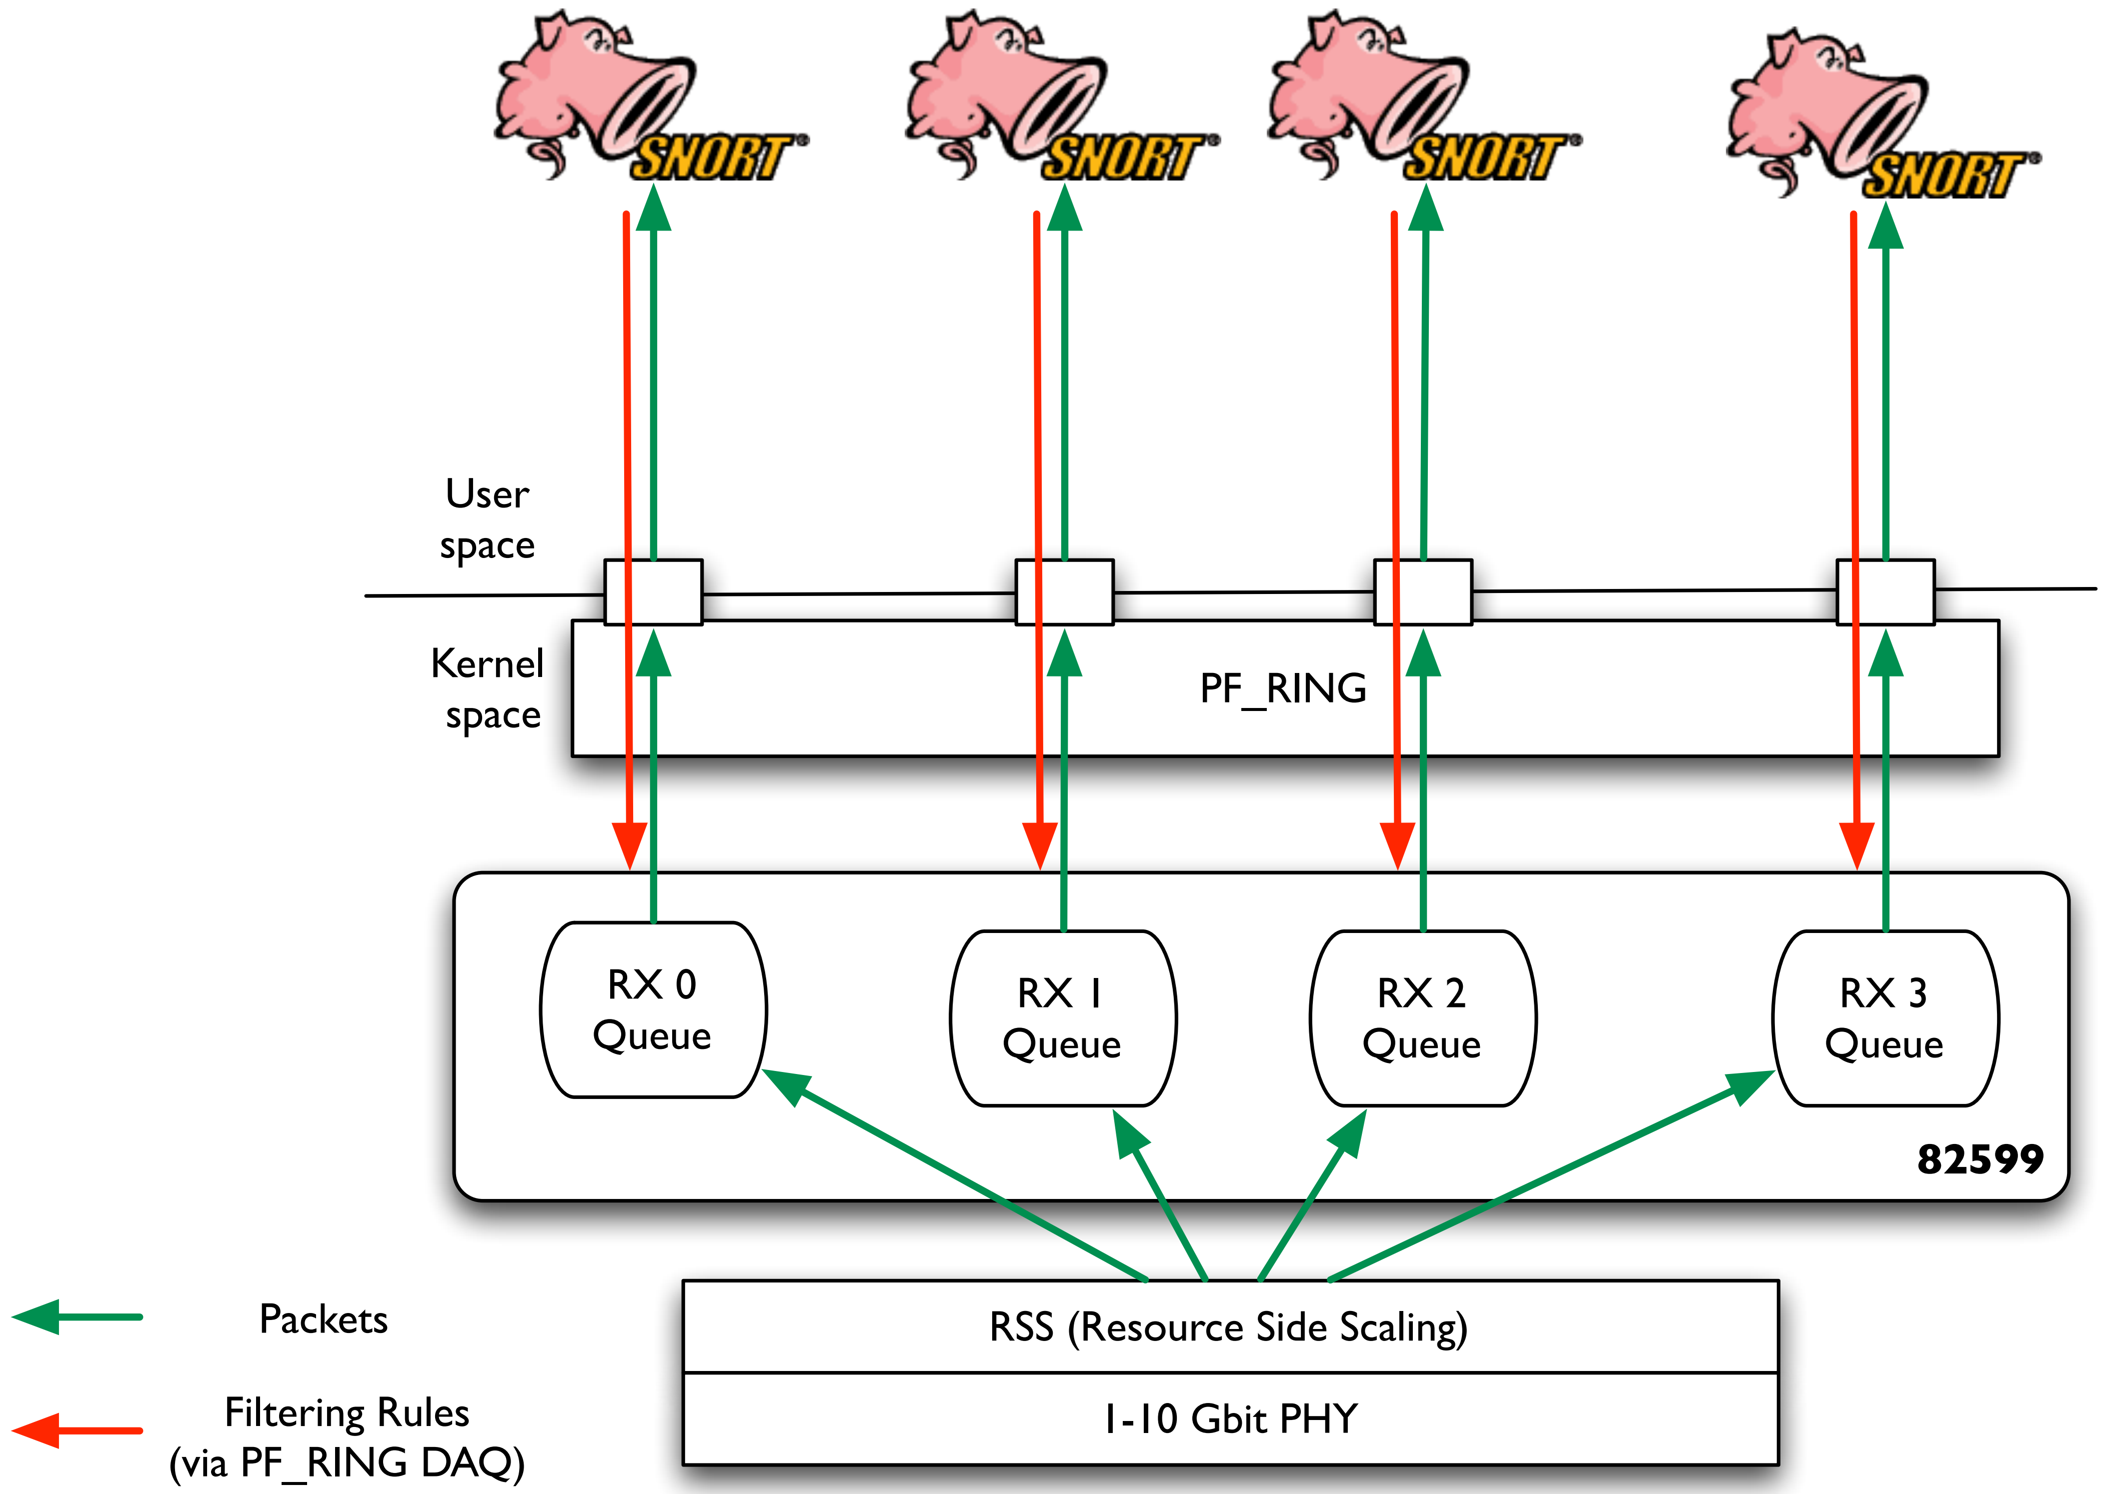

User space

Kernel space

PF_RING

RX 0 Queue

RX 1 Queue

RX 2 Queue

RX 3 Queue

82599

RSS (Resource Side Scaling)

1-10 Gbit PHY

Packets

Filtering Rules (via PF_RING DAQ)

Divide et Impera

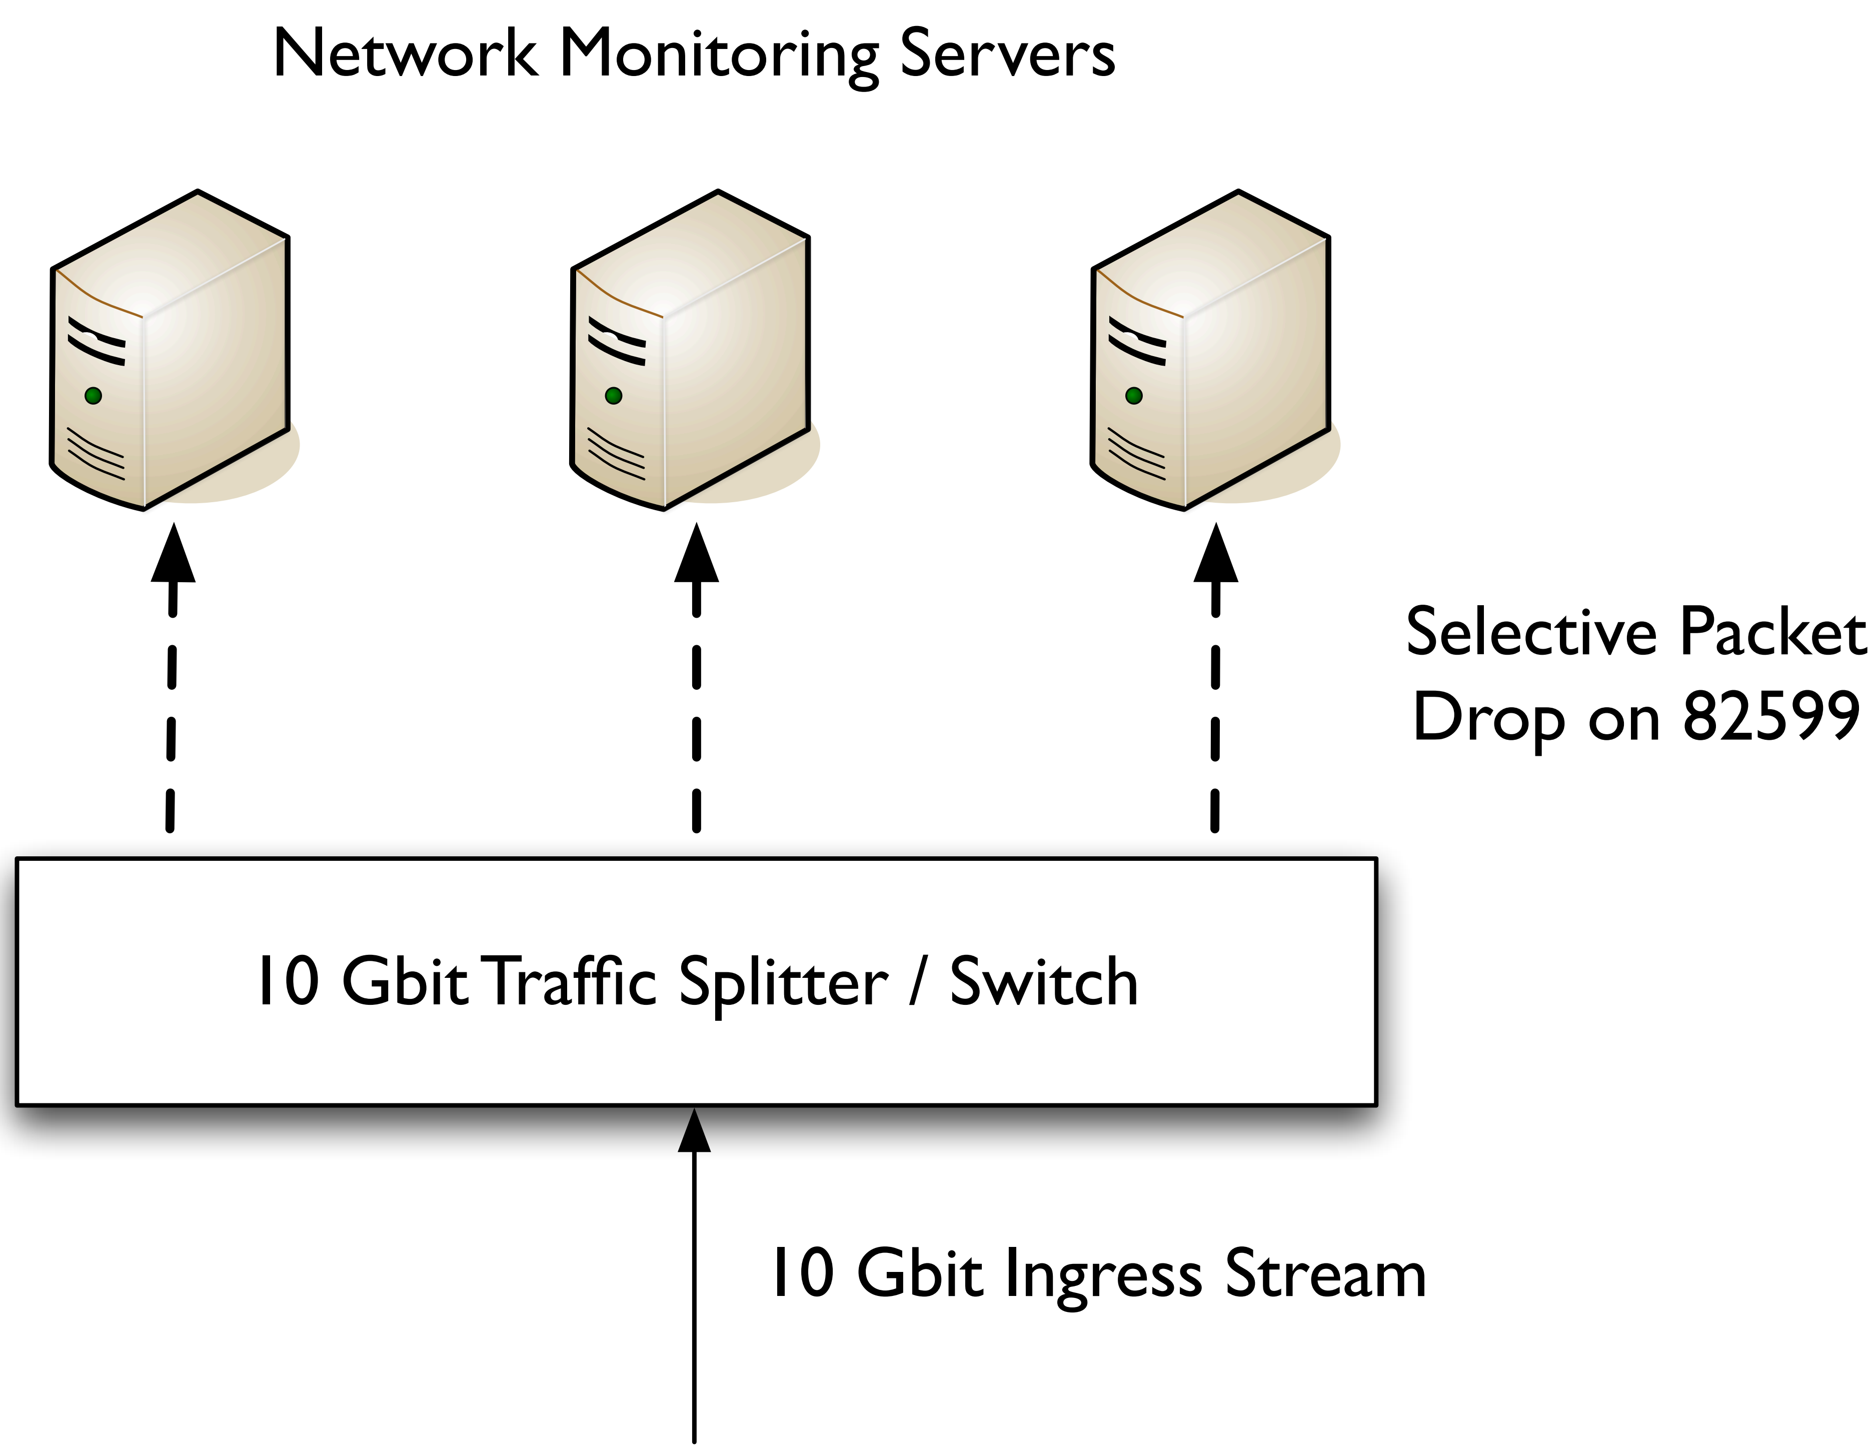

Network Monitoring Servers

Selective Packet Drop on 82599

10 Gbit Traffic Splitter / Switch

10 Gbit Ingress Stream

Performance Figures [1/2]

Single filtering rule matching all packets

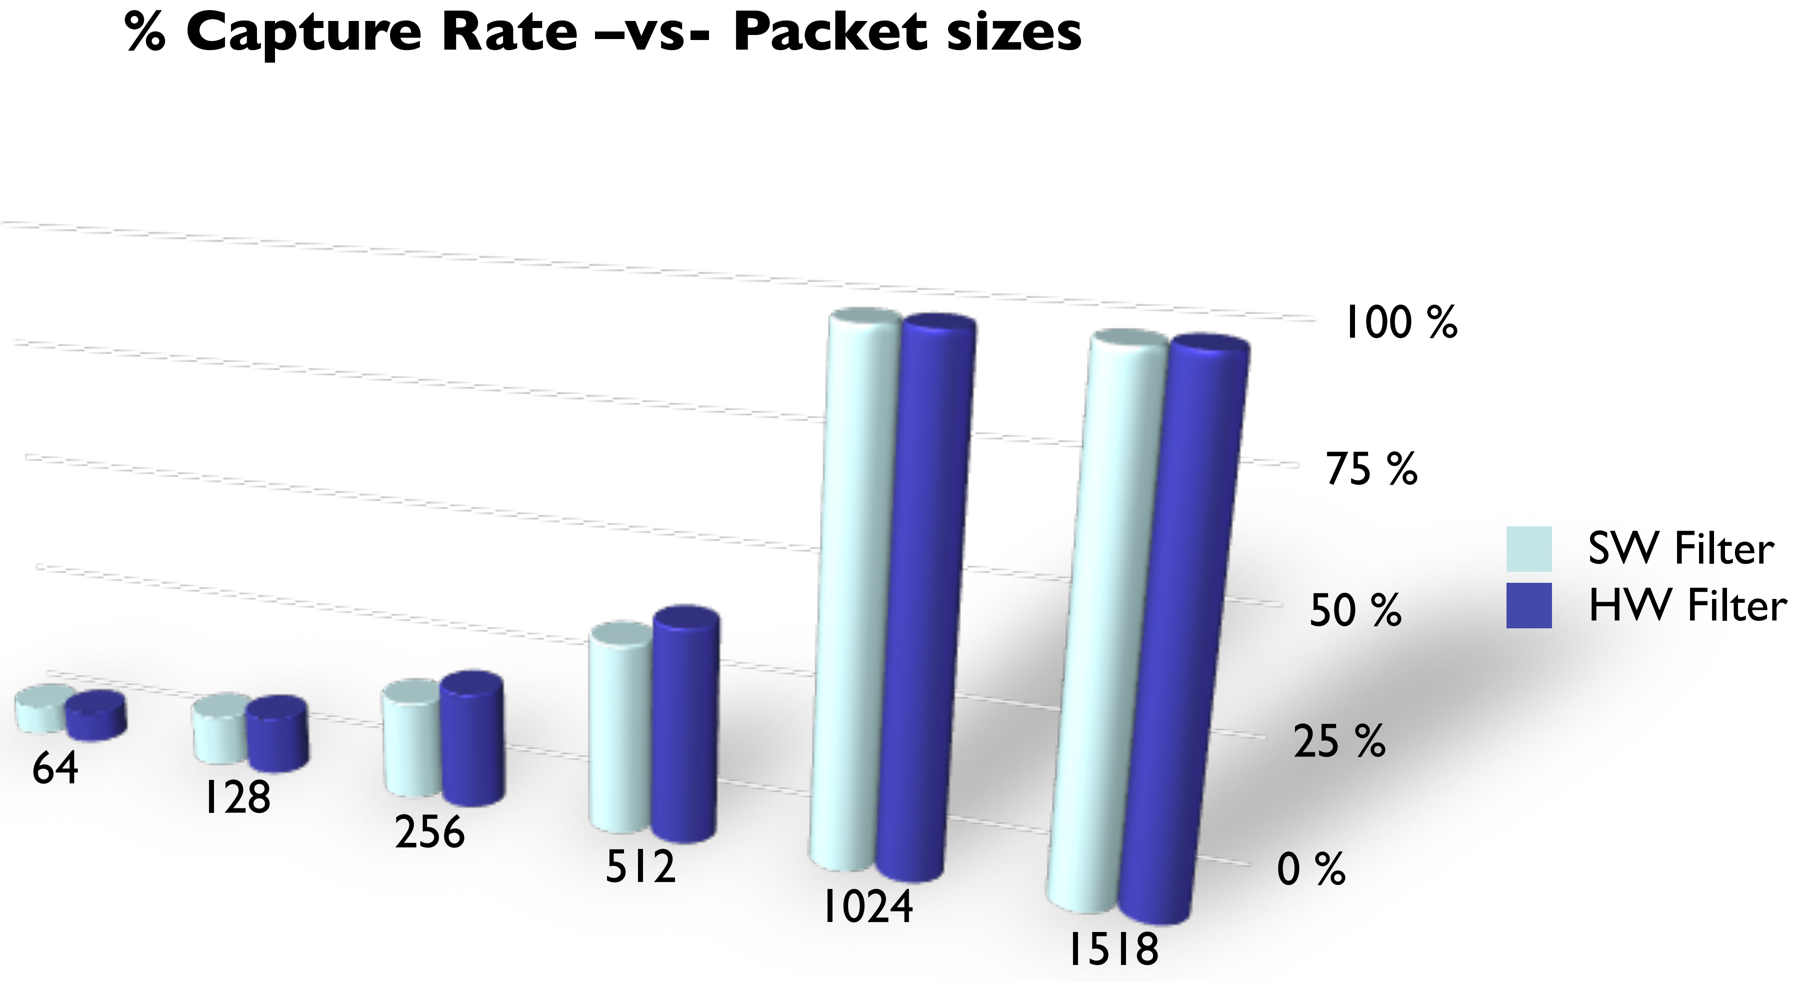

Note: Using multiple filters increases significantly CPU usage when SW filters are used (Butterfly Effect), whereas filter number does not affect HW filters.

Performance Figures [2/2]

Single filtering rule dropping all packets

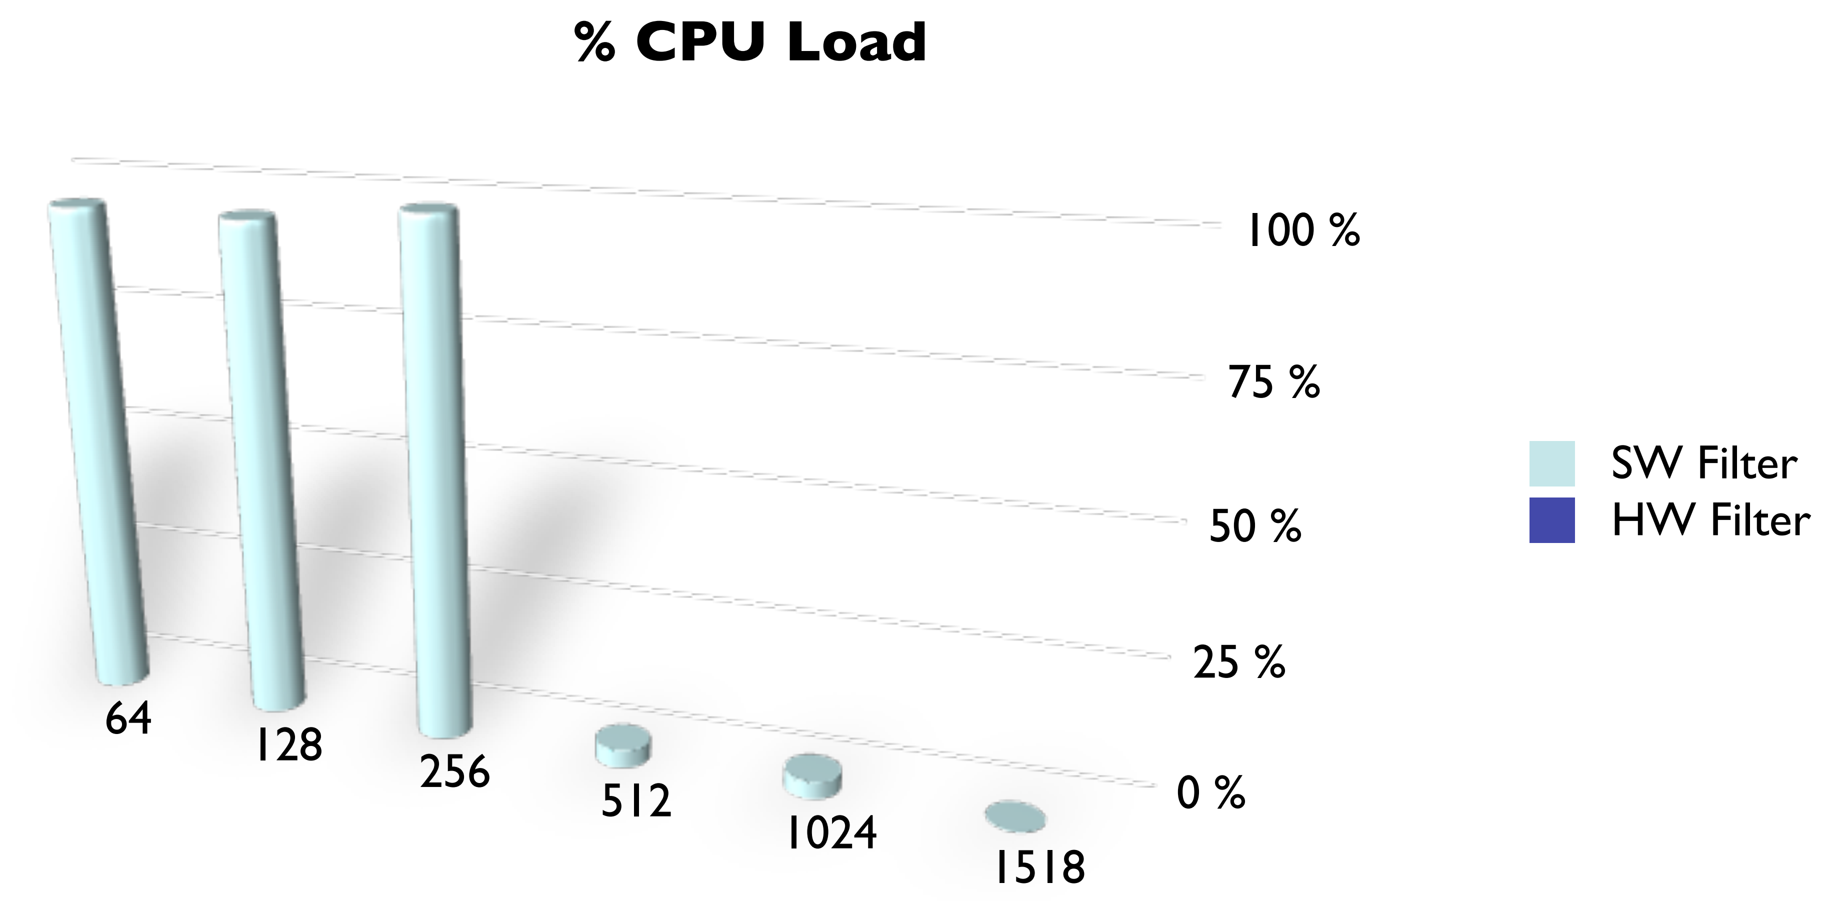

Note: As expected CPU is not loaded at all when HW filters are used.

Final Remarks

- Using hardware-assisted packet filtering and balancing allows network administrators to monitor and troubleshoot 10 Gbit networks using commodity hardware.

- Available at no cost (GNU GPL) from http://www.ntop.org/PF_RING.html

- L. Deri, J. Gasparakis and F. Fusco
Wire-Speed Hardware Assisted Traffic Filtering with Mainstream Adapters
Proceedings of NEMA 2010, October 2010