

# Towards 10 Gbit NetFlow Monitoring Using Commodity Hardware

Luca Deri <deri@{unipi.it,ntop.org}>

# 10G Technology Overview

- 10Gbit is available in various PHY (6 for fiber, 3 for copper), the most popular/cheap is 10GBASE-SR (fiber 850nm)
- Retention of 802.3 MAC and frame format
- Different from other versions of Ethernet
  - No half duplex mode
  - 10G only: no 10/100/1000/10G
- Works with 802.1Q, 802.3ad, etc.
- 10 GE is still an emerging technology with only 1 million ports shipped in 2007.
- PC adapters prices are falling (< 1000 Euro), PCI-X adapters replaced by PCIe.

# 10 Gbit NetFlow Challenges

- High number of packets to be analyzed (10 times as much as 1 Gbit).
- CPU-based traffic analysis (e.g. as it happens in most router-based netflow probes) is not feasible at these speeds: dedicated cards are needed.
- Packet filtering is very important, in particular on WANs, in order to early discard those packets that are supposed not to be analyzed.
- Support of WAN encapsulations (e.g. MPLS)

# nProbe Overview

- Support for IPv4/v6 and NetFlow v5/v9/IPFIX.
- Ability to act as a NetFlow/IPFIX probe, proxy (protocol converter) and collector.
- Ability to operate at 1 Gbit wire speed on commodity hardware.
- Support for major OS (Unix, Windows and MacOS X) as well strong multicore systems (Tilera Tile64)
- Resource (both CPU and memory) savvy, efficient, designed for environments with limited resources.
- Source code available under GNU GPL.

# Scaling to 10 Gbit: Divide et Impera

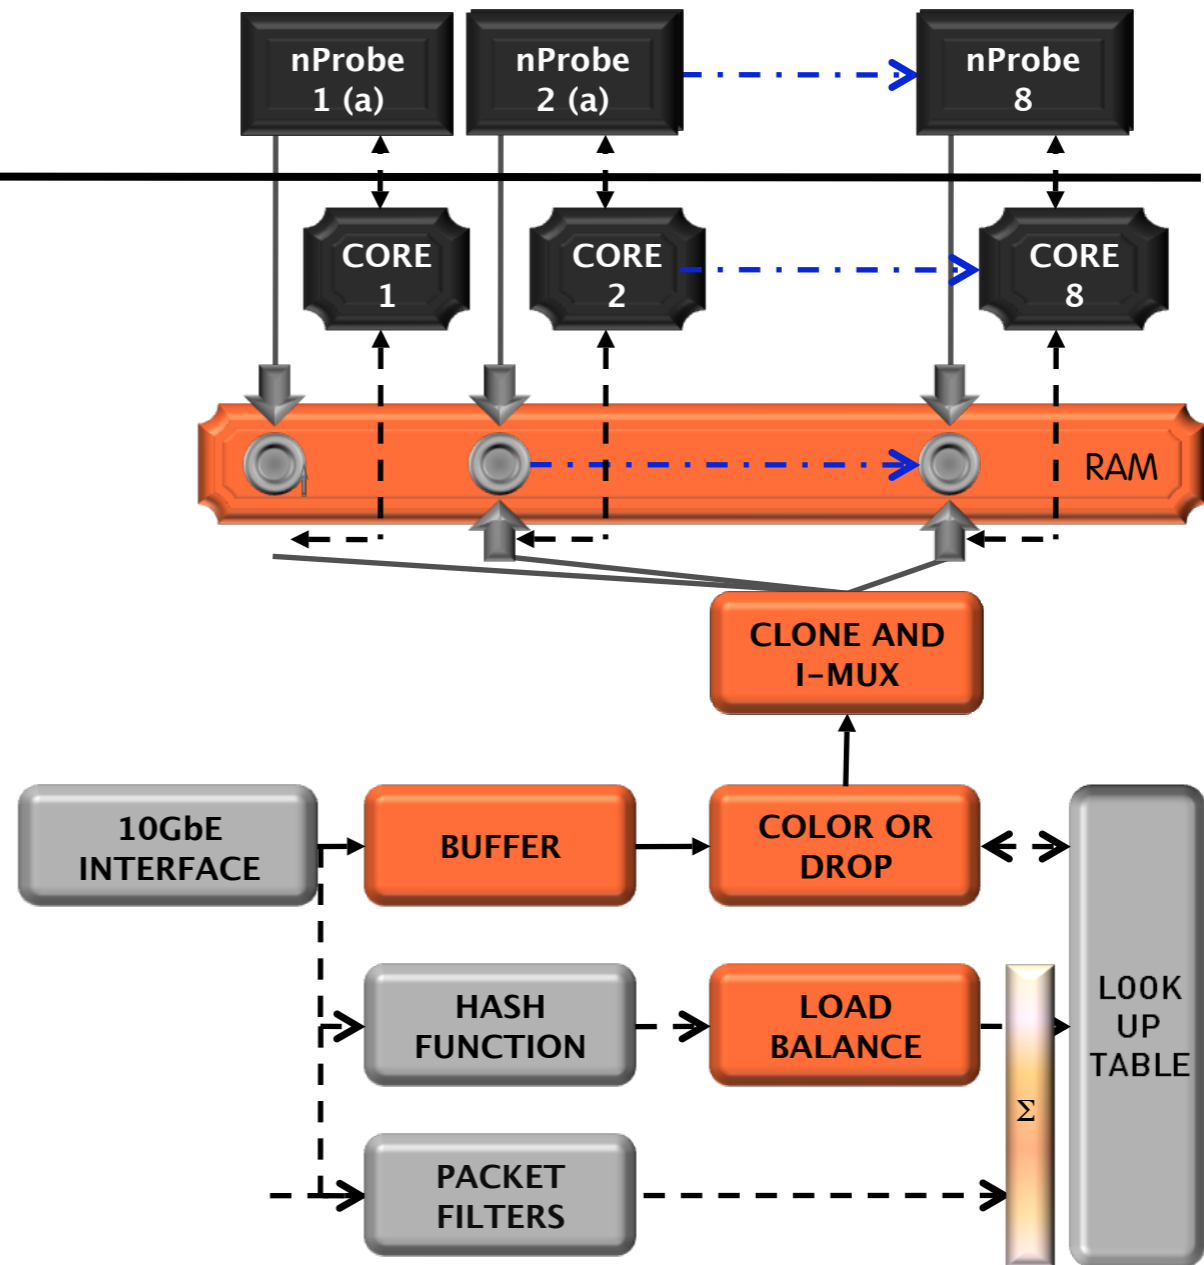
- CPU manufacturer are scaling with multicore.
- Multicore equations:
  - more cores = **more total** CPU power
  - more cores = **less single** core power
- Software scales with multicore only if it can exploit it:
  - multiprocess or multithread
- A “simply faster” 10 Gbit NIC is not enough:
  - one 10G card means that several threads need to compete for packets hence that a lot of time will be wasted on mutexes

# nProbe on a 10 Gbit DAG

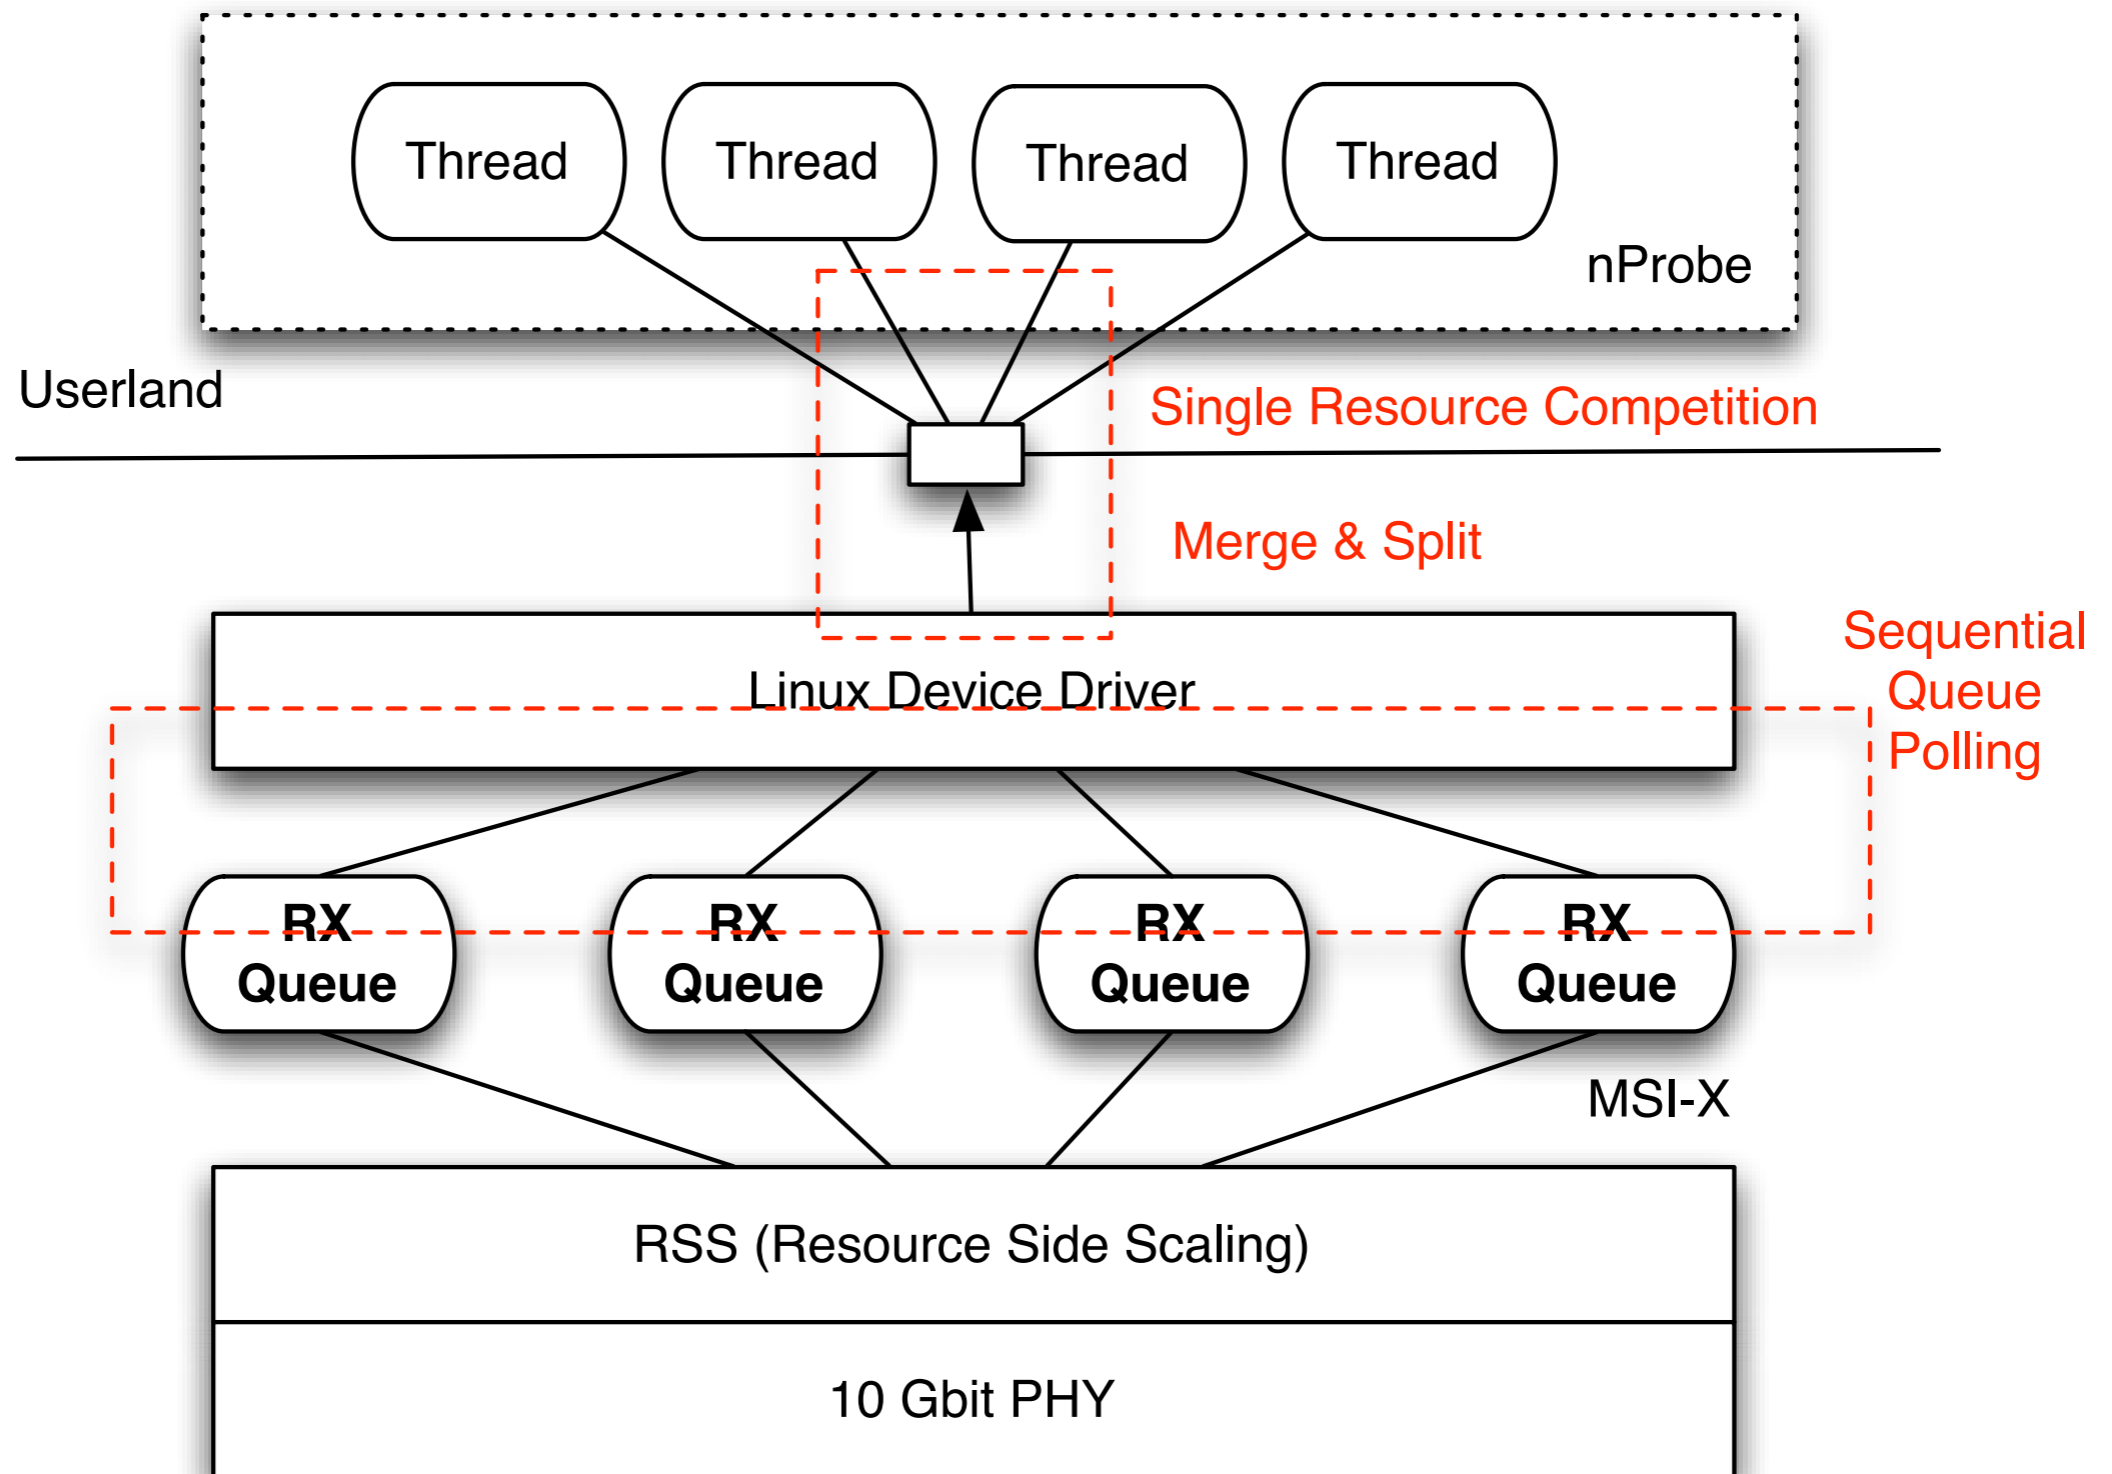
Userland

Kernel

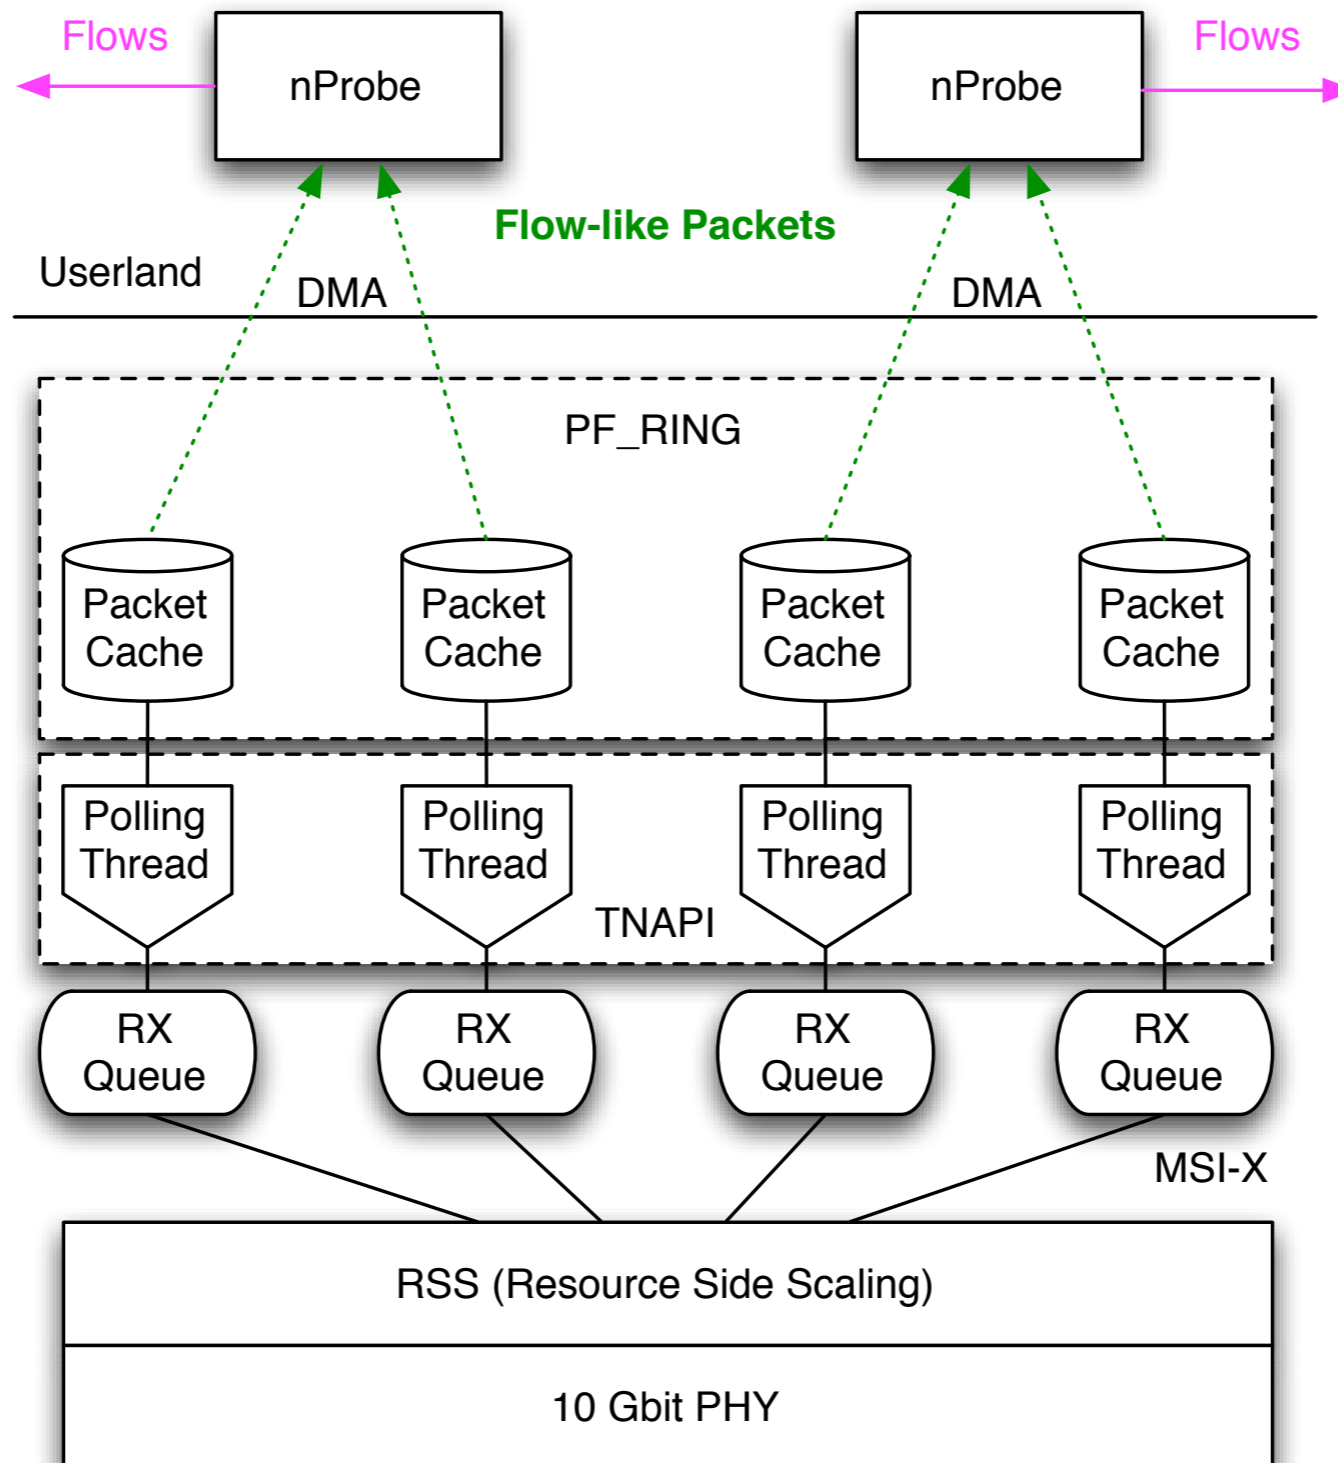
- Packet balancing across 8 nProbes/cores.
- Peak nProbe performance: 100% Packet capture and flow processing up to ~6 Mpps with no sampling.



# Multicore+Networking Design Flaws

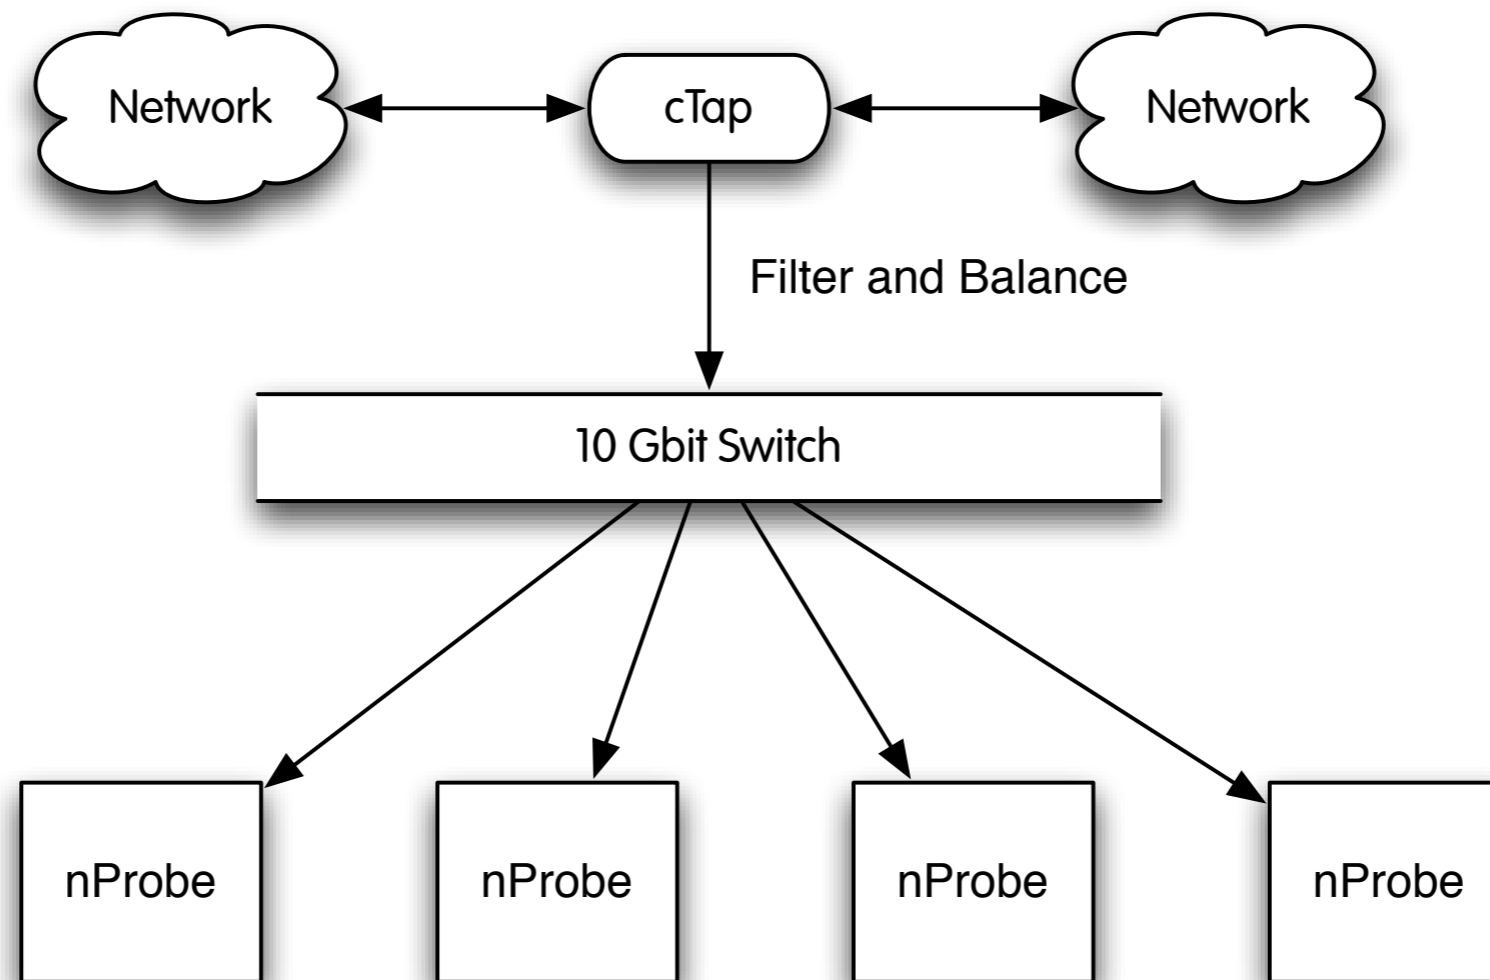


# nProbe + PF\_RING + TNAPI



- Packet balancing across cores.
- Peak nProbe performance: 1.48 Mpps (packet rate) x 2 Cores.

# Further Divide et Impera



# Summary

- 10 Gbit NetFlow monitoring (not just packet capture) is possible using open source probes such as nProbe.
- The basic assumption is that the monitored traffic must be 'balanceable'.
- Scaling to 40 and 100 Gbit is also possible with multiple 'divide et impera' iterations.
- The new challenges are now on the collector side: will it be able to handle all the monitoring data?