

---

# Hardware-Accelerated NetFlow Probe

Ladislav Lhotka  
*<lhotka@cesnet.cz>*

Martin Žádník  
*<zadnik@liberouter.org>*



- Data about IP flows proved to be useful for a number of purposes:
  - ▷ Traffic statistics and accounting
  - ▷ Network planning
  - ▷ Security analysis
  - ▷ QoS monitoring
- Cisco NetFlow is the de facto standard
- Routers are the primary source (ASIC for NetFlow)
- Variety of tools available (FTAS, NERD, NFSen, Stager, ...)

- Not all routers and switches support NetFlow
- Routers as L3 devices are vulnerable to attacks
- Routers have limited capacity (CPU, TCAM)
- Multiple purposes often result in conflicting requirements (sampling, aggregation, ...)
- Routers cannot do additional processing (anonymisation, per-collector filtering)
- Little room for testing new features (advanced sampling techniques, user-defined templates, ...)

# FlowMon probe

---



- Linux PC with HW accelerator card acting as a GE repeater – stealth device
- Simultaneous monitoring of IPv4 and IPv6 traffic
- Hardware-accelerated header parsing, matching of key fields and statistics gathering
- Format-independent flow records, export in NetFlow v5 and v9, per-collector filtering
- Cache for 64 Kflows
- Standard statistical sampling, sample-and-hold


# FlowMon hardware

---



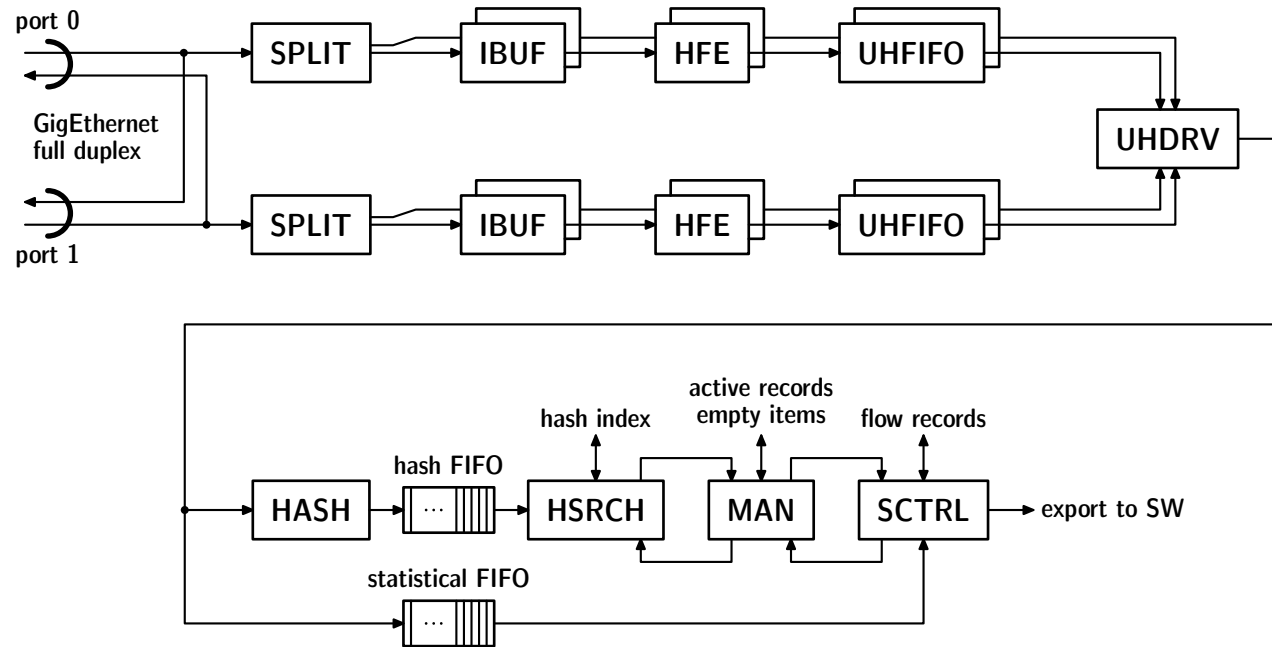
- COMBO cards with Xilinx Virtex II-Pro FPGAs
  - ▷ COMBO6X – PCI-X motherboard
  - ▷ COMBO-4SFPRO – 4× Gigabit Ethernet (SFP)
- Previous version used COMBO6 (32/33 PCI) and COMBO-4SFP
- COMBO-2XFP – 10GE card, must be redesigned due to unavailability of components
- Other uses of COMBO cards: Liberouter, SCAMPI, IDS, packet generator, DNA sequencing

# FlowMon hardware (cont.)

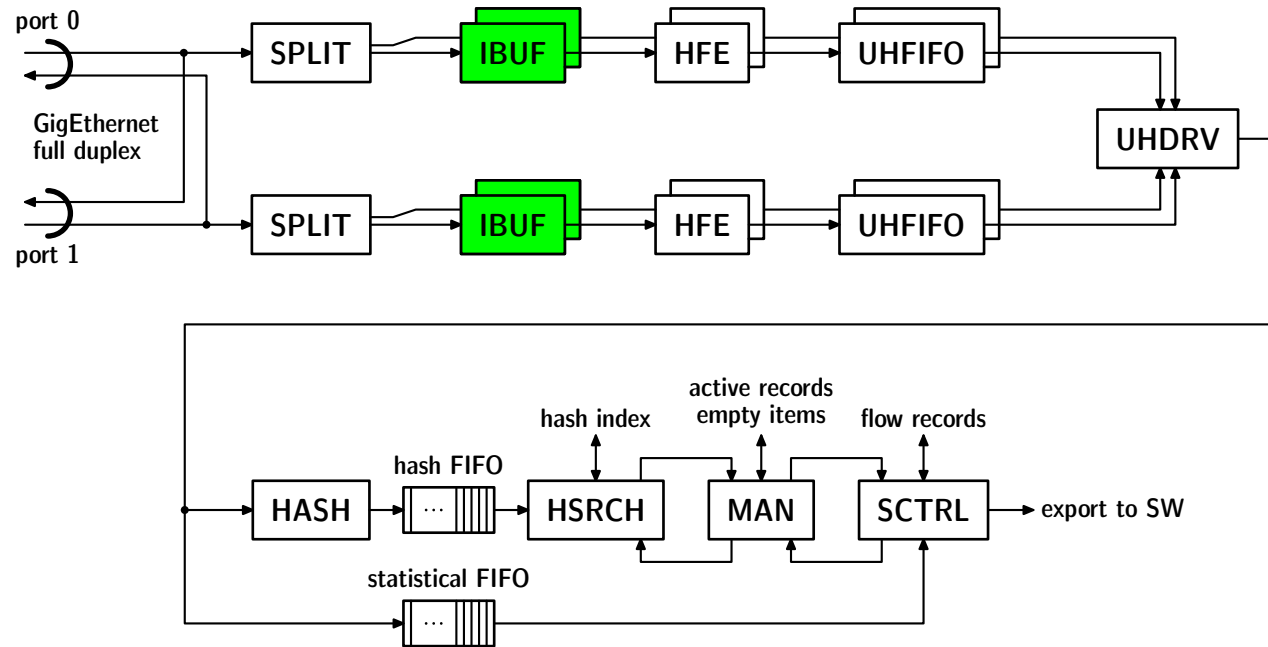


COMBO6X (bottom) and COMBO-4SFPRO (top)

# FlowMon firmware

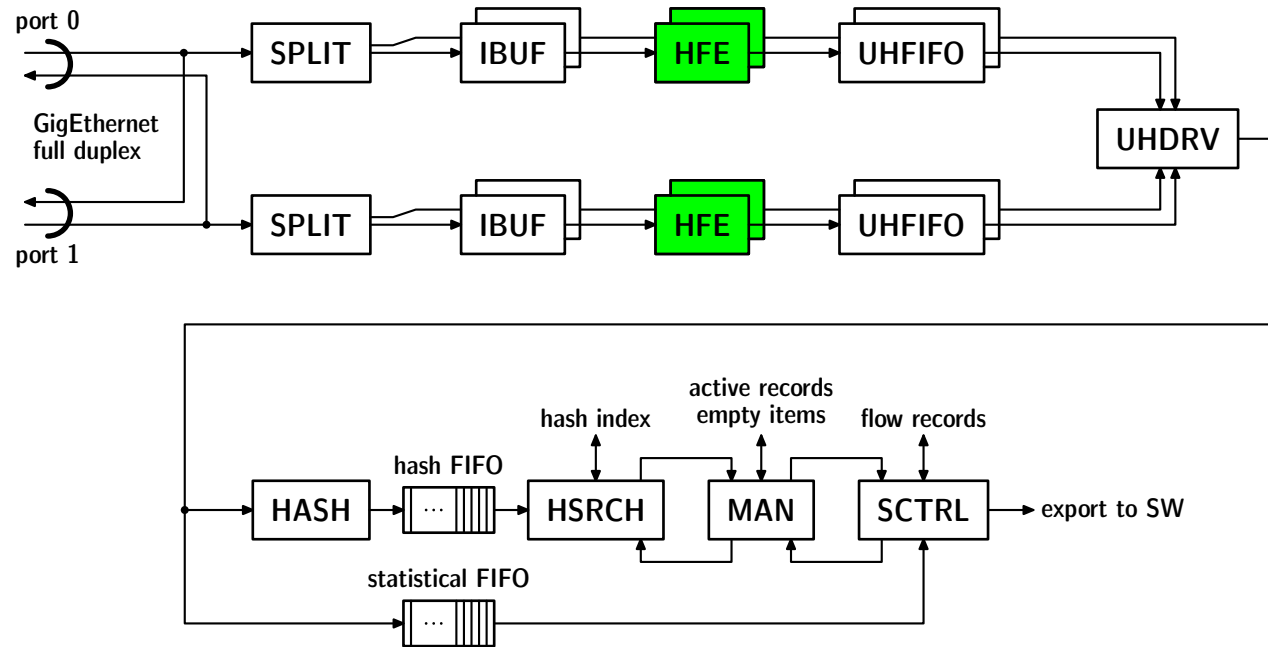


# FlowMon firmware (cont.)



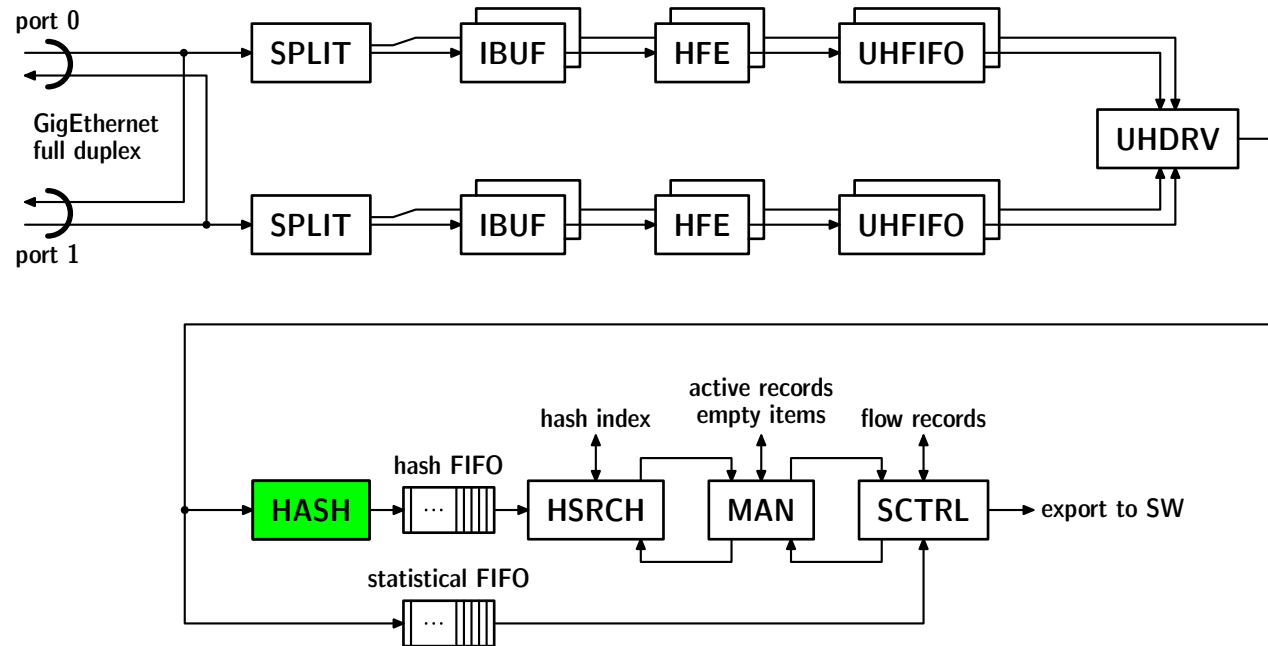
**Input Buffer** also implements statistical sampling.

# FlowMon firmware (cont.)



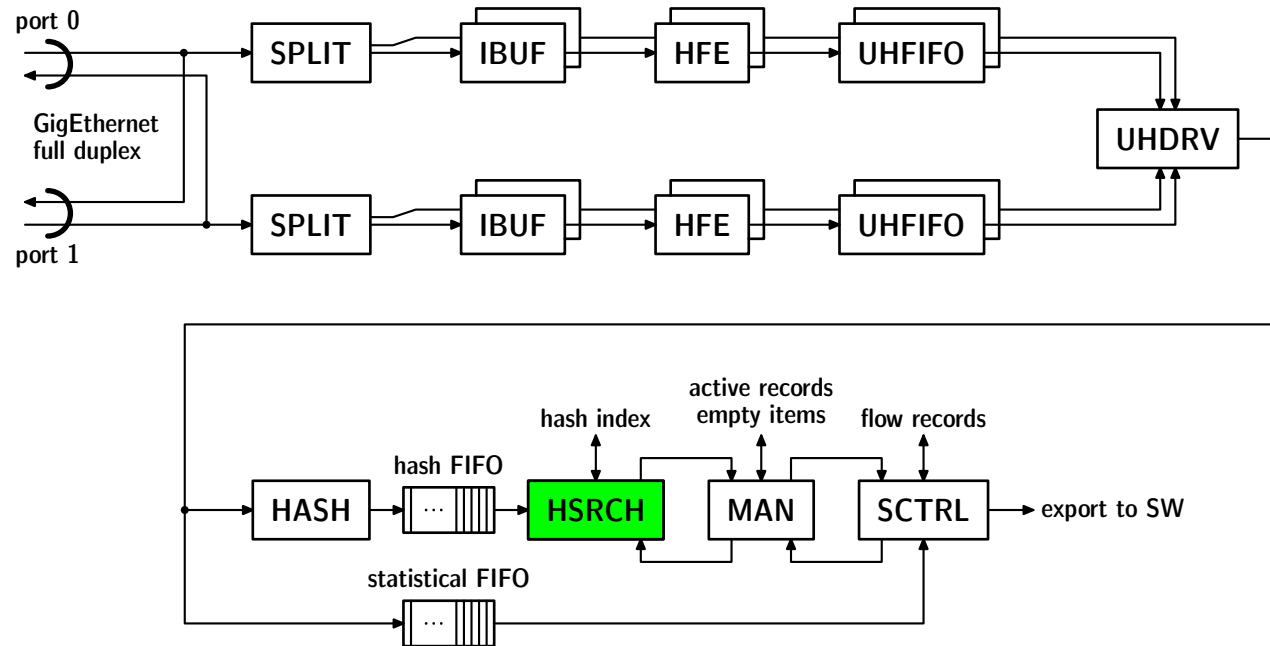
**Header Field Extractor** parses packet headers and collects relevant fields into *Unified Header (UH)*.

# FlowMon firmware (cont.)



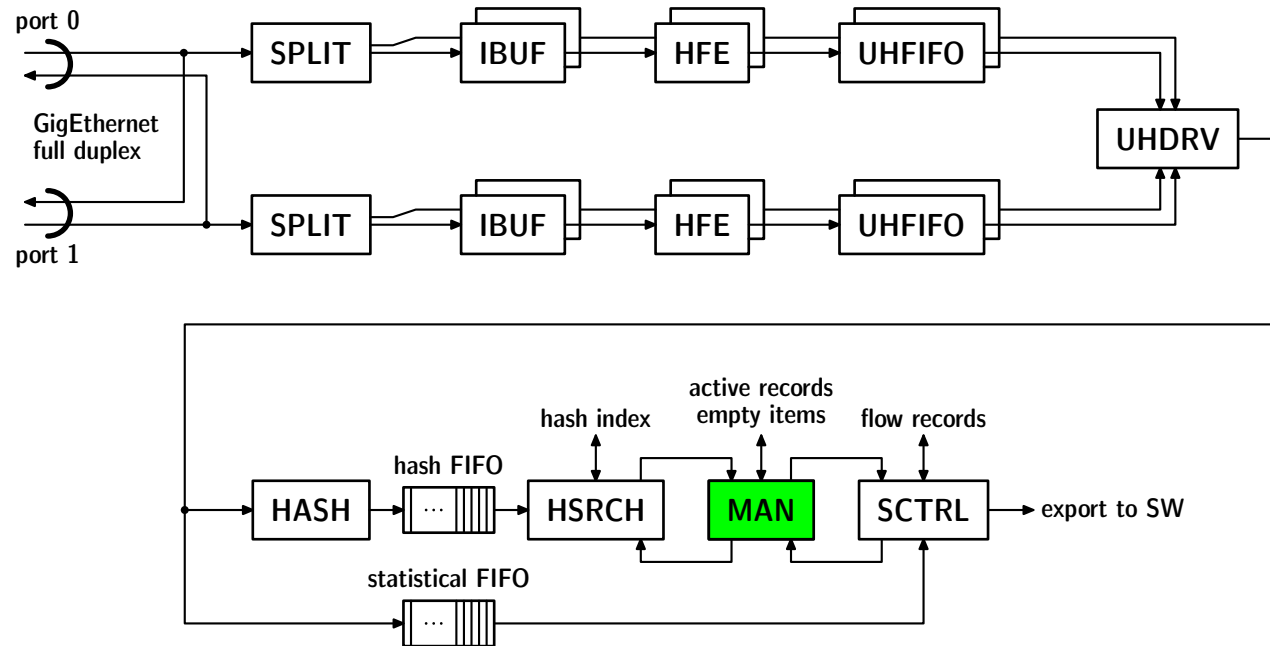
**Hash Unit** computes CRC-64 hash from predefined UH *key fields* (IP addresses, ports, protocol).

# FlowMon firmware (cont.)



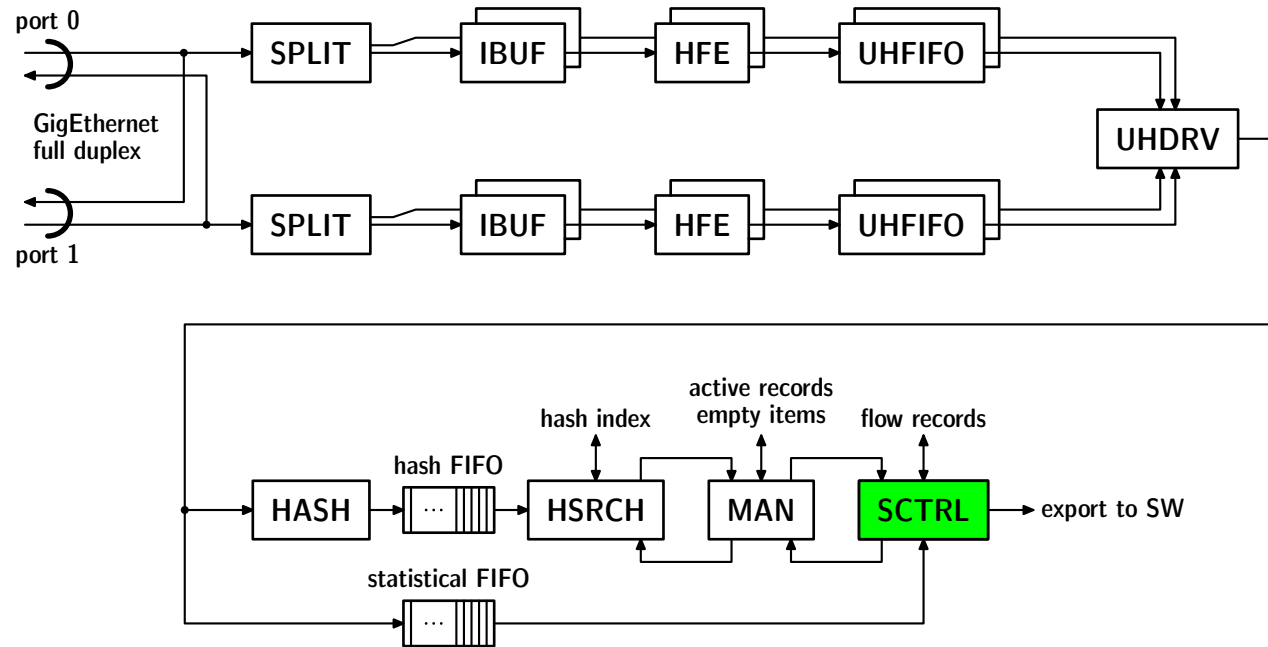
**Hash Search Unit** tries to find a matching hash in the index and reports the result to MAN, implements sample-and-hold.

# FlowMon firmware (cont.)



**Manager Unit** controls bookkeeping of active flows and free memory locations, expires flows according to timeouts.

# FlowMon firmware (cont.)



**Storage Control Unit** manages flow records, updates statistics, exports expired flow records.

# Hash function

---



- Key fields are too long, esp. for IPv6
- We use 57 bits of their CRC-64 value
- Collision probability  $\leq 2^{-41} \doteq 4,55 \times 10^{-13}$  (less than 15 collisions a year at 1 Mpps on the average)
- To avoid provoked collisions, HASH is initialised with a random seed.

# Device driver

---



- Several cooperating Linux kernel modules (2.4 and 2.6 kernels supported)
- Flow records are stored in a shared memory block
  - circular buffer for 16K flow records (new records rewrite oldest ones)
- Multiple applications may access the data simultaneously, each with its own read pointer
- Each application may lock up to 1024 records

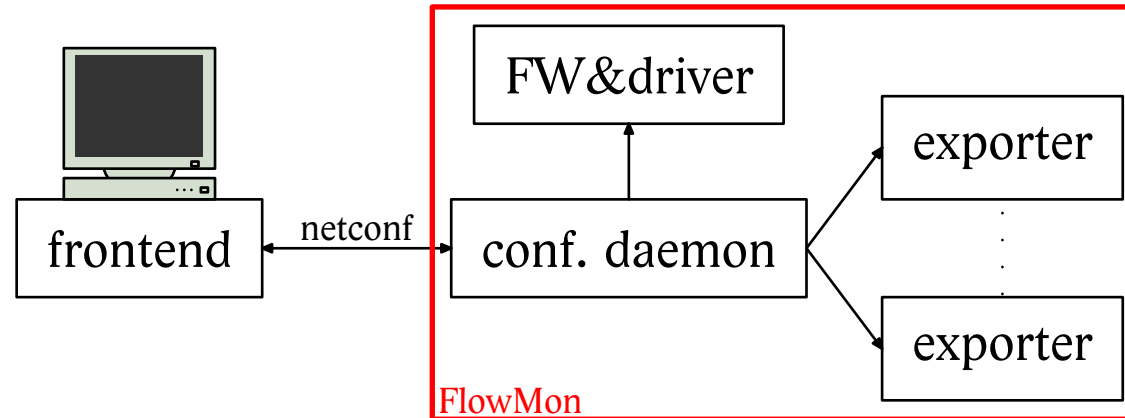
# Flow export programs

---



- NetFlow v5
- NetFlow v9 (RFC 3954)
  - ▷ 8 templates (TCP, UDP, ICMP, OTHER for both IPv4 and IPv6)
  - ▷ Templates sent periodically (configurable)
  - ▷ Options not supported yet
- Both exporters support multiple collectors
- Per-collector filtering of records based on source and/or destination address ranges

# Configuration



- Text frontend generates XML configuration and passes it to the configuration daemon *flowmond* (via disk file, netconf planned).
- Other frontends planned (graphical, web)
- *flowmond* configures firmware, driver and one or more exporter programs

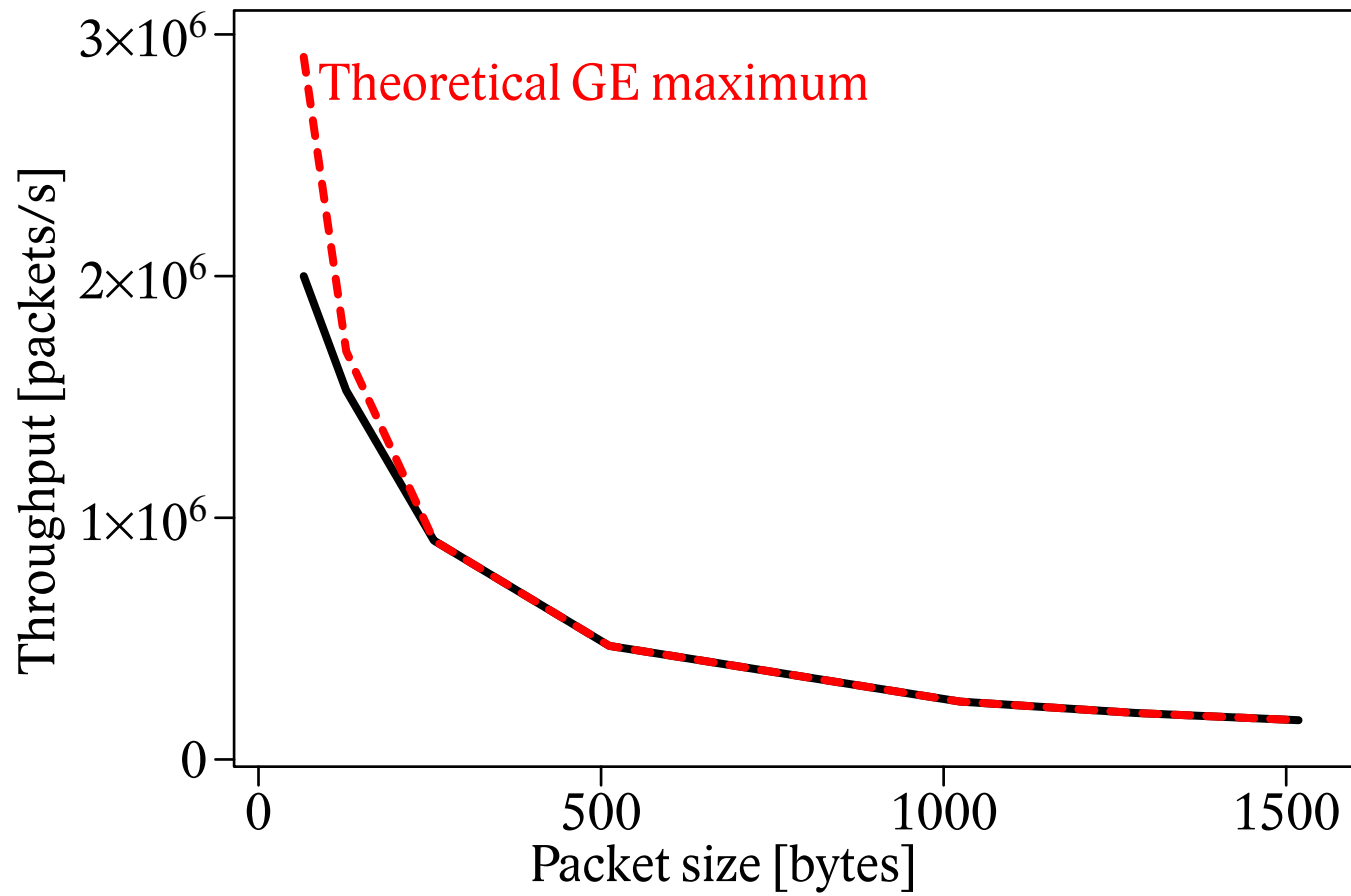
# Configuration options

---



- Firmware parameters
  - ▷ Active and inactive timeout
  - ▷ Sampling rate, sample-and-hold option
  - ▷ Logging options
- For each configured collector:
  - ▷ hostname/address and port
  - ▷ export protocol (NetFlow v5 or v9)
  - ▷ flow record filter (multiple address ranges to be applied to SrcIP, DstIP or either of them)

# Full duplex performance



# Field tests

---



- Two probes deployed in CESNET backbone, more planned
- CESNET lent four probes to JRA2 partners (SURFnet, SWITCH, GRNET, Bulgarian Acad. Sci.)
- Successfully tested with NERD, NfSen, FTAS, for IPv4 and IPv6, NetFlow v5 and v9

# Future plans

---



- Hardware
  - ▷ 10GE interface card
  - ▷ SDH STM-16
- Firmware
  - ▷ throughput 3 Mpackets/s
  - ▷ cache for 500 Kflows
  - ▷ configurable flow record
  - ▷ anonymisation

# Future plans (cont.)

---



- Software
  - ▷ IPFIX
  - ▷ user-defined templates
- Research topics
  - ▷ adaptive sampling
  - ▷ additional HW functions (aggregation, heuristics)
- Commercial availability and support

# Summary

---



- FlowMon probe is being developed as part of security monitoring toolset of JRA2@GN2
- Flexible platform for both routine monitoring and research, integrates well with the existing collectors and other tools
- NetFlow v5 and v9 supported, IPFIX planned
- Capable of processing full-duplex GE (later SDH STM-16) at line rate without sampling, 10 Gb/s with sampling

- Home page

*<http://www.flowmon.org/flowmon-probe>*

- GN2 deliverable DJ2.2.2

*<http://www.geant2.net/server/show/nav.00d00b002>*