



Modern Network Traffic Monitoring

Visibility, Security, Awareness

Luca Deri
deri@ntop.org, @lucaderi

Introduction

- This talk is about creating a comprehensive, high-speed traffic monitoring platform using commodity hardware and (open source) software components.
- Price, complexity, limited vendor support (bytes 'n packets), and proprietary (i.e. not fully standard compliant) implementations have made people perceive network monitoring as just “nice to have”.
- We demonstrate that it's now possible to effectively monitor traffic with limited effort leveraging on recent computing advances.
- Most tools presented in this talk are available on github at <https://github.com/ntop>.
 -  full open source
 -  some open source components

Packets, Flows, Activities [1/3]

- For years monitoring tools focused on standards often fostered by vendors: NetFlow vs sFlow vs SNMP, Cisco vs Juniper...
- This has plagued the market by creating tools more vendor- oriented, than result- oriented.
- Fortunately recent advances in computing and in particular the big data movement, have pushed companies to overcome the market/vendor fragmentation and produce tools able to produce data on a standard format (often JSON) that could be consumed even by non-monitoring tools (e.g. Hadoop, ElasticSearch).

Packets, Flows, Activities [2/3]

- As data increases and people demand feature rich monitoring tools, it has become necessary *to compress* monitoring data.
- Network packets are still important for providing evidence or troubleshooting problems (packets or it didn't happen!) but they are “too raw” and take too much storage space, so limiting them to specific situations is a good idea.
- Network flow analysis is a good way to “compress packets” into events: sFlow do it with sampling, NetFlow with stateful connection-based packet classification.

Packets, Flows, Activities [3/3]

- These days, saving flows on a big data system is a common practice but it still plagued by the visibility issue:
 - What flows are “more relevant” than others?
 - Can we use flows for more than just host/protocol/application traffic accounting ?
 - How can a network administrator look for a needle in a haystack when the monitoring platform is emitting tenth of thousand flows/second?
- We need yet another level of abstraction on top of flows able to identify activities on top of flows (e.g. these 20 HTTPS connections and 5 DNS queries mean that host X just open the landing page of newspaper corriere.it).

Flow Generation [1/2]

- Unfortunately there are still too many “NetFlow dialects” (e.g. Cisco ASA or Barracuda Networks flows) available that make interoperability not that simple
- sFlow is even simpler than NetFlow/IPFIX to implement and available in most switches deployed today (Cisco features a sFlow-like protocol named NetFlow Lite).
- With the baseline bytes/packets of traffic flow from these flow protocols, we can do a lot with good analytics. This including congestion, cost analysis, DDoS detection, security and forensics.

Flow Generation [2/2]

- Ideally, we want to gather *rich measurement metrics*, from everywhere possible.
- For the above goals 5-tuples (IPs, Ports, Protocol) and bytes/packets are not enough as we expect at least:
 - Latency, Packet Drops, Retransmissions.
 - QoE (e.g. HTTP service time).
 - Application visibility (DPI, URLs, DNS responses).
- And with those metrics per flow, we can provide even *more actionable insights* into performance and security issues.

Affordable High-Speed Everywhere [1/3]

- “From packets” vs. “from flow” can yield the enriched metadata that is useful for modern operations. Sensors that see packets from taps or SPAN ports are now affordable enough to deploy widely.



Affordable High-Speed Everywhere [2/3]

- Caveat: sensors must be *affordable* and *rich* in measurement metrics (we need more than bytes/packets).
- What is high-speed today?
 - 1 Gbit for home and small offices.
 - 10 Gbit for medium business.
 - 40/100 Gbit for ISP/large business
- In order to make network monitoring commodity, *all the traffic* must be monitored, *not just the core network*.

Affordable High-Speed Everywhere [3/3]

- What is a price ballpark for a line-rate sensor (hardware and software) that could be placed everywhere on a network?
 - 1 Gbit: 1'000 Euro
 - 2 x10 Gbit: 2'500 Euro
 - 100 Gbit: 25'000 Euro
- Anything more than this, isn't considered *affordable* and thus *everywhere*. Later in this presentation you will find out how to make this possible in 2016.


100 Gbit Packet Capture [1/2]

- FPGA-based NICs (e.g. Accolade and Napatech) have been out for more than a year now, with prices (not including optics) starting in the 10K USD range.
- Recently Intel has introduced the FMI0000 10/25/40/100 Gbit Ethernet controller (Red Rock Canyon) that offloads to the controller selected features (e.g. packet switching/distribution/drop). The first products have been announced and are available since 1Q16 for < 1.5k USD (dual 100 Gbit).



100 Gbit Packet Capture [2/2]

- For monitoring full-duplex 100 Gbit links using network taps, it is necessary to:
 - Spread packets across NUMA nodes.
 - Respect flow coherency (each core sees both traffic directions).
 - Vendors overcome this problem using various technologies.



← Napatech Custom Interconnection Cable

From 10 Gbit to 100 Gbit [1/3]

- Just as RSS (Resource Side Scaling) allows network adapters to distribute traffic across cores on modern network adapters, it is possible to do the same using network switches.
- Some products (e.g. Dell Z9100-ON) allow traffic to be statically (no RSS-like features) distributed across multiple 10/25/40 Gbit ports for reducing 100 Gbit monitoring to multi 10-Gbit monitoring.



From 10 Gbit to 100 Gbit [2/3]

- OpenFlow-programmable switches, can be very well used to distribute ingress traffic across their interfaces and this implement the “divide and conquer” paradigm.
- You can setup ACLs (same as FMI000) for selectively dropping/filtering/diverting traffic across ports.
- ACLs do not support any sort of RSS-like features, meaning that you cannot dynamically balance traffic across ports, but you need to implement this statically.

From 10 Gbit to 100 Gbit [3/3]

- In essence both Red Rock Canyon and external OpenFlow-programmable switches can achieve the same goal with a different form-factor.
 - ACLs can help dropping/filtering/redirecting traffic but they are not dynamic and flow-aware.
 - These solutions are good for selected cases where the traffic to analyse is very specific and predictable in terms of IPs and ports.
 - For all other cases, a beefy server with a 100 Gbit adapter is the best option.

Flows for Everything, pcap for Something [1/4]

- Many organisations need to:
 - Have evidence of *all* activities happened in their network (generally less true for ISPs who have different needs).
 - Satisfy this requirement using packet recorders, that can store packets to disk in pcap format.
- As traffic rate increase, this approach is no longer working and new solutions need to be identified:
 - 10 Gbit: 1.25 GB/sec
 - 40 Gbit: 5 GB/sec
 - 100 Gbit: 12.5 GB/sec

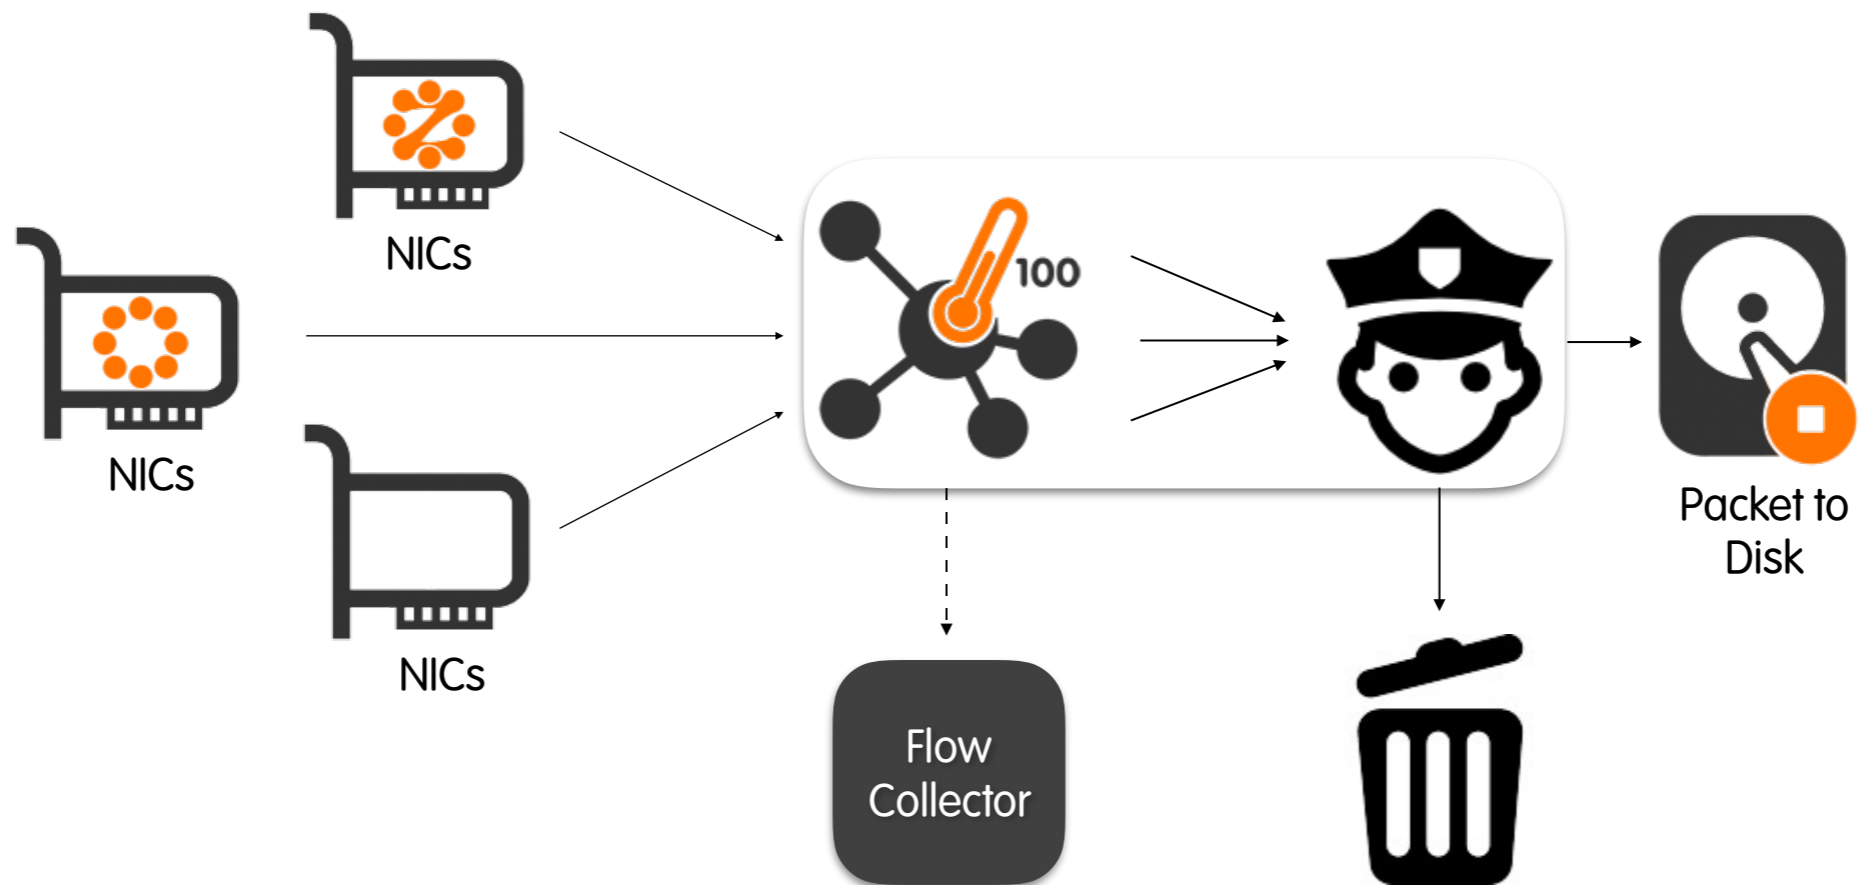
Flows for Everything, pcap for Something [2/4]

- Storage space is not the only reason why not all traffic has to be recorded to disk:
 - Encrypted traffic can be of little help in case of attacks.
 - Multimedia streams (e.g. Netflix or AppleMusic) can take significant unnecessary disk space.
- Traffic compression does not help on the Internet as most traffic is already compressed (e.g. JPEG, HTML).
- Space is not just a cost, as it affects speed: searching for packets on a smaller pcap puts less pressure on the storage system (and thus avoid packet drops).

Flows for Everything, pcap for Something [3/4]

- Generally, recording *all traffic* to disk *is not* a good idea. Instead recording all the *interesting* traffic to disk *is* a good idea.
- So our current approach is to generate enhanced flow from *everything* and store *specific* (for example unencrypted, non-video) packets.
- Flows must be generated on all the traffic as we need have to keep evidence of all activities, whereas packets need to be generated only on relevant events as saving 40G/100G to disk is simply unfeasible (if useful at all).

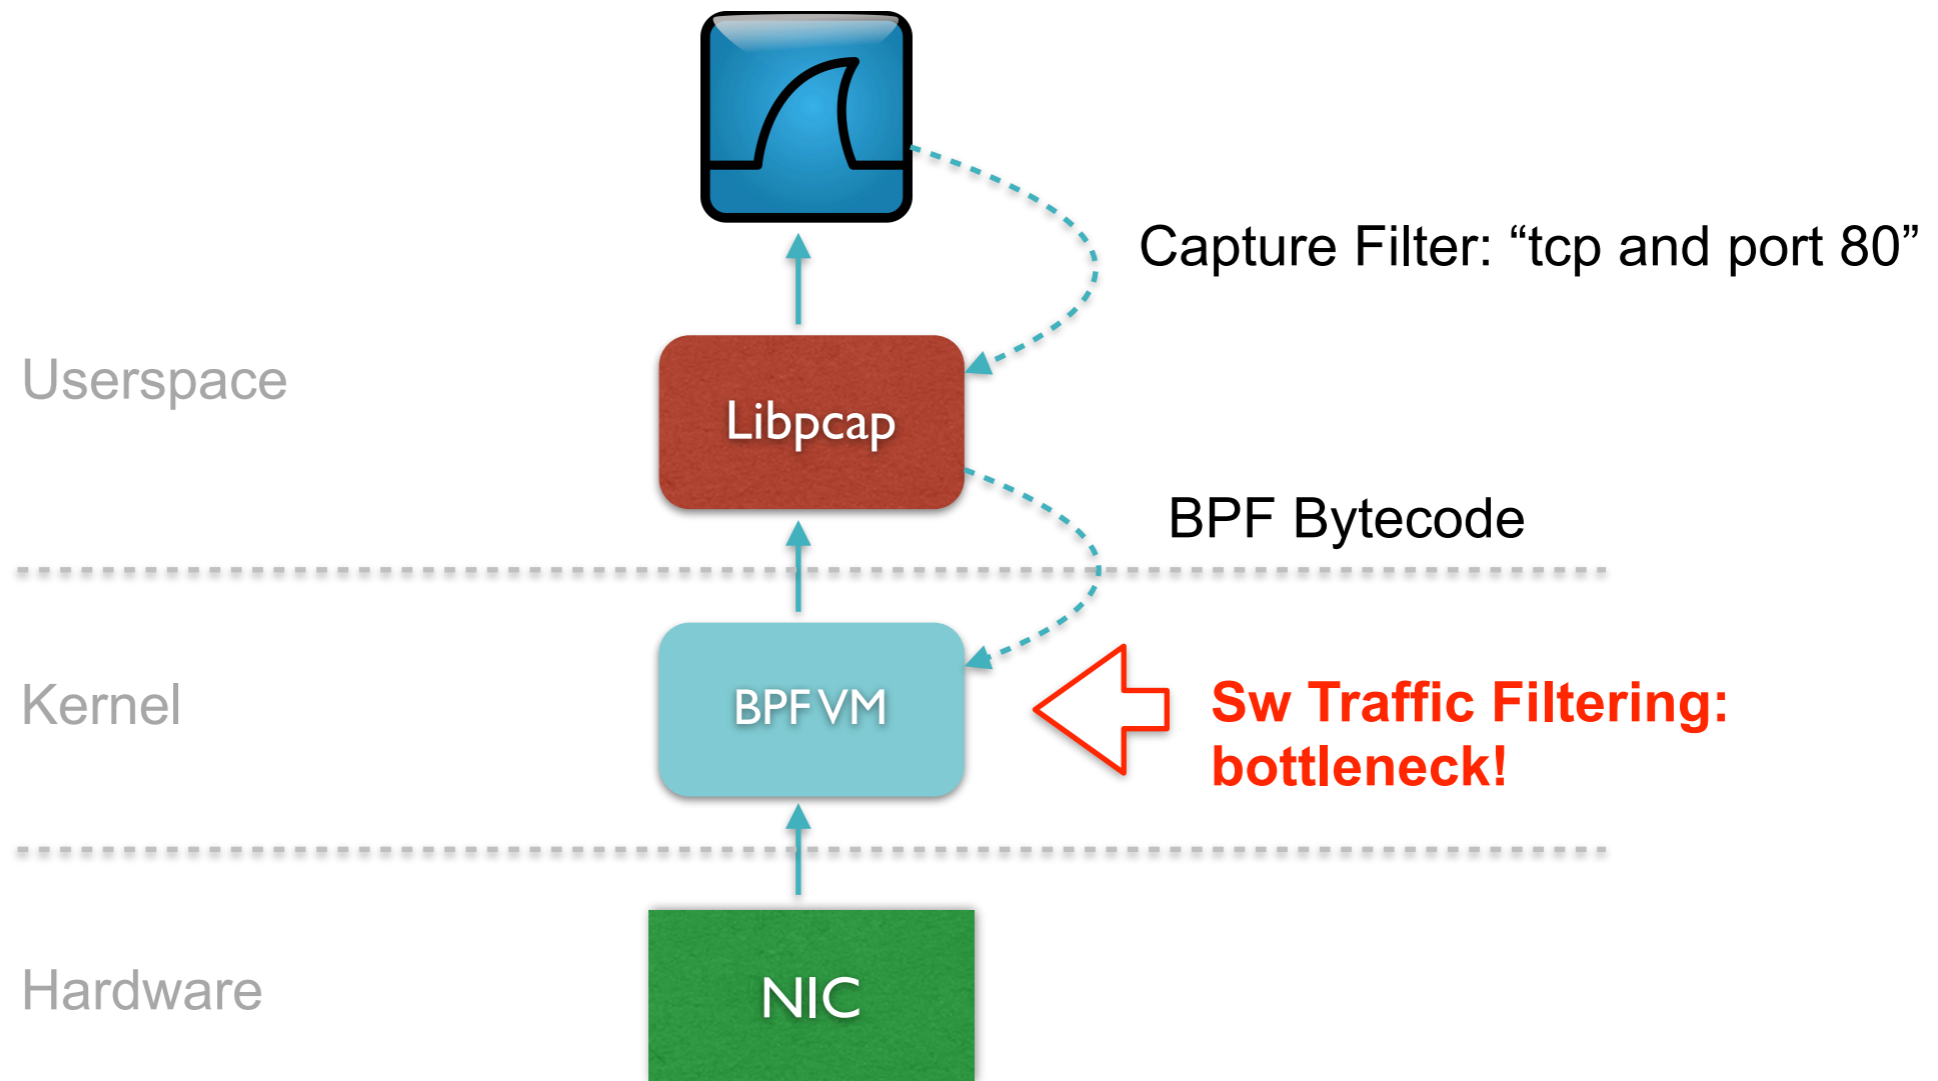
Flows for Everything, pcap for Something [4/4]



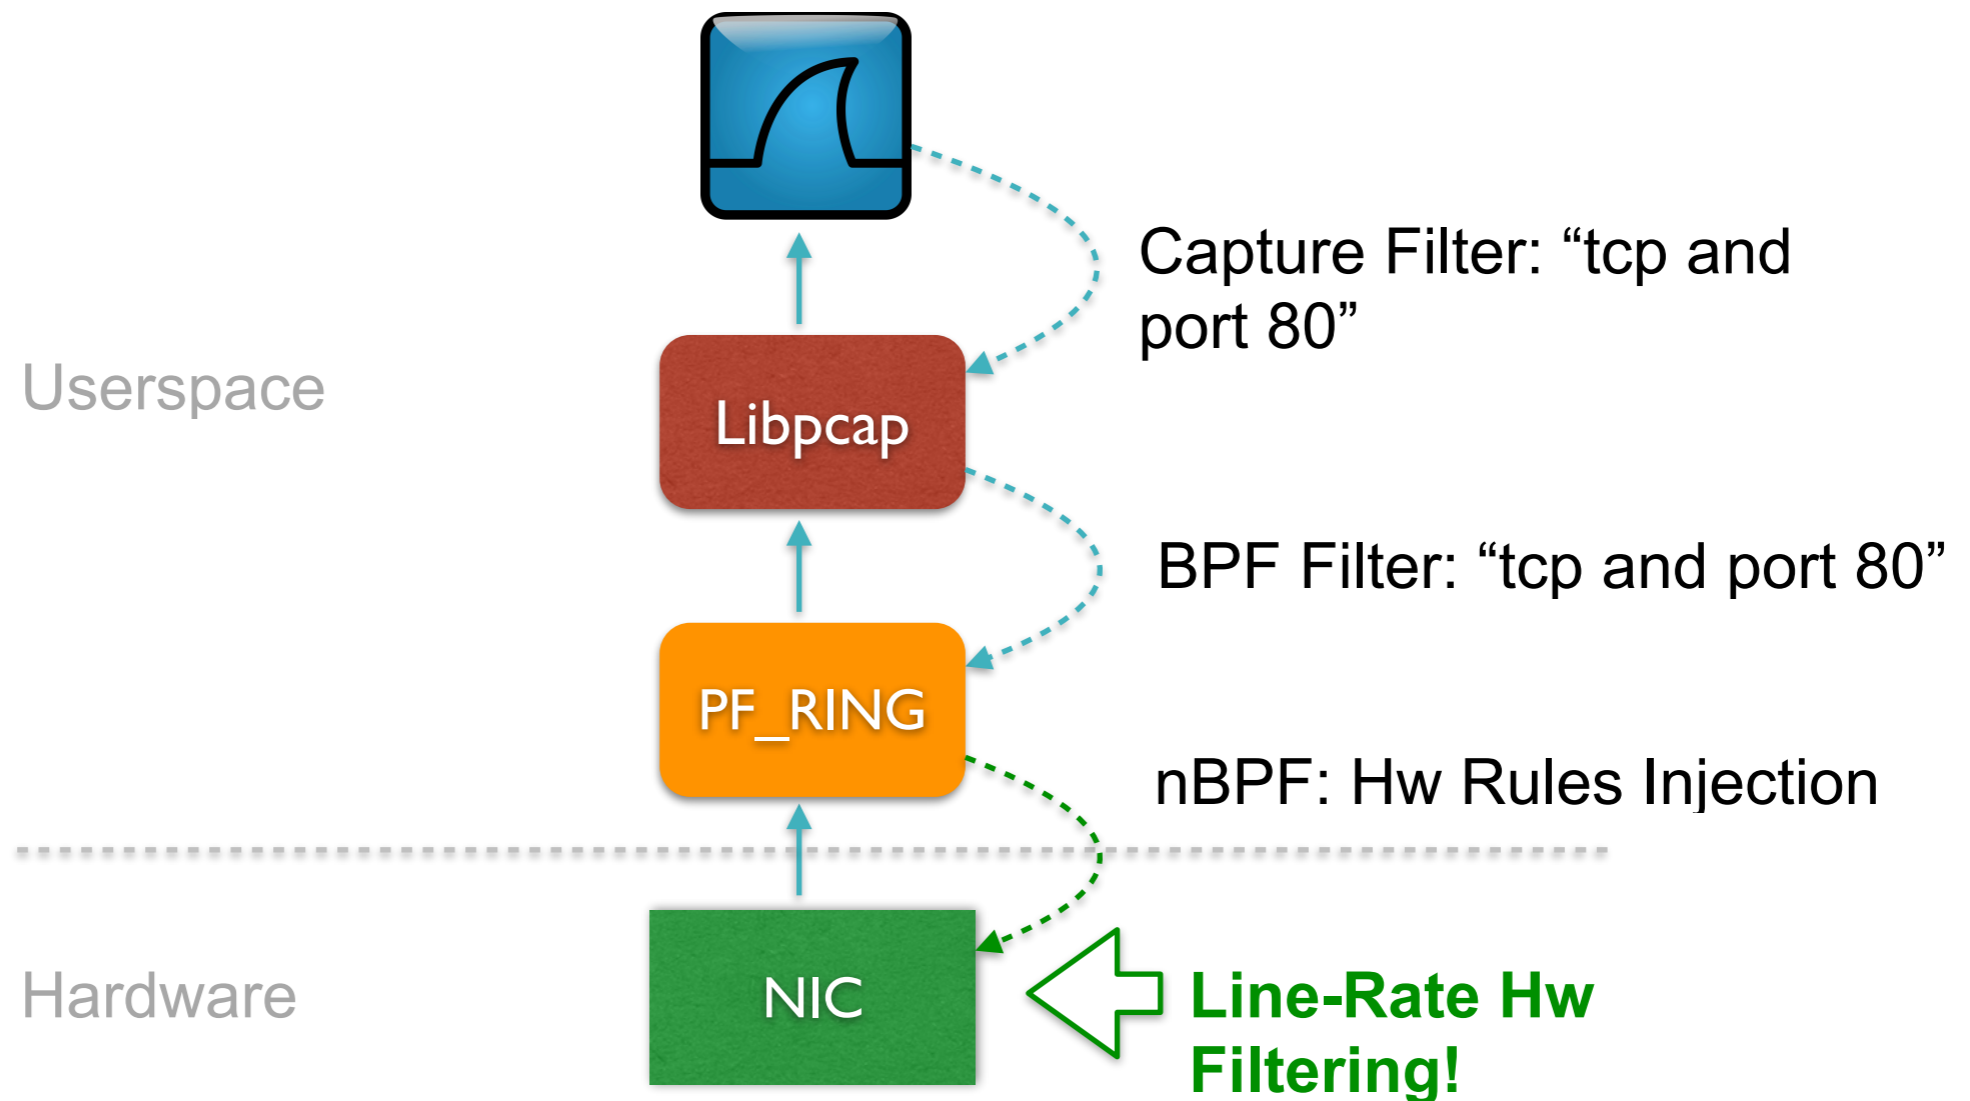
Troubleshooting at 100 Gbit [1/3]

- Sometimes a “quick & dirty” debug session it is enough to solve a network problem.
- Starting Wireshark on a 40/100 Gbit link to capture live traffic is not a good idea unless you know exactly what you are looking for (e.g. host X and port Y).
- Even in this case the traffic is far too much on a busy link thus traffic filtering is compulsory.

Troubleshooting at 100 Gbit [2/3]



Troubleshooting at 100 Gbit [3/3]



 ntop's nBPF is available for Napatech, Intel FM10K, Exablaze

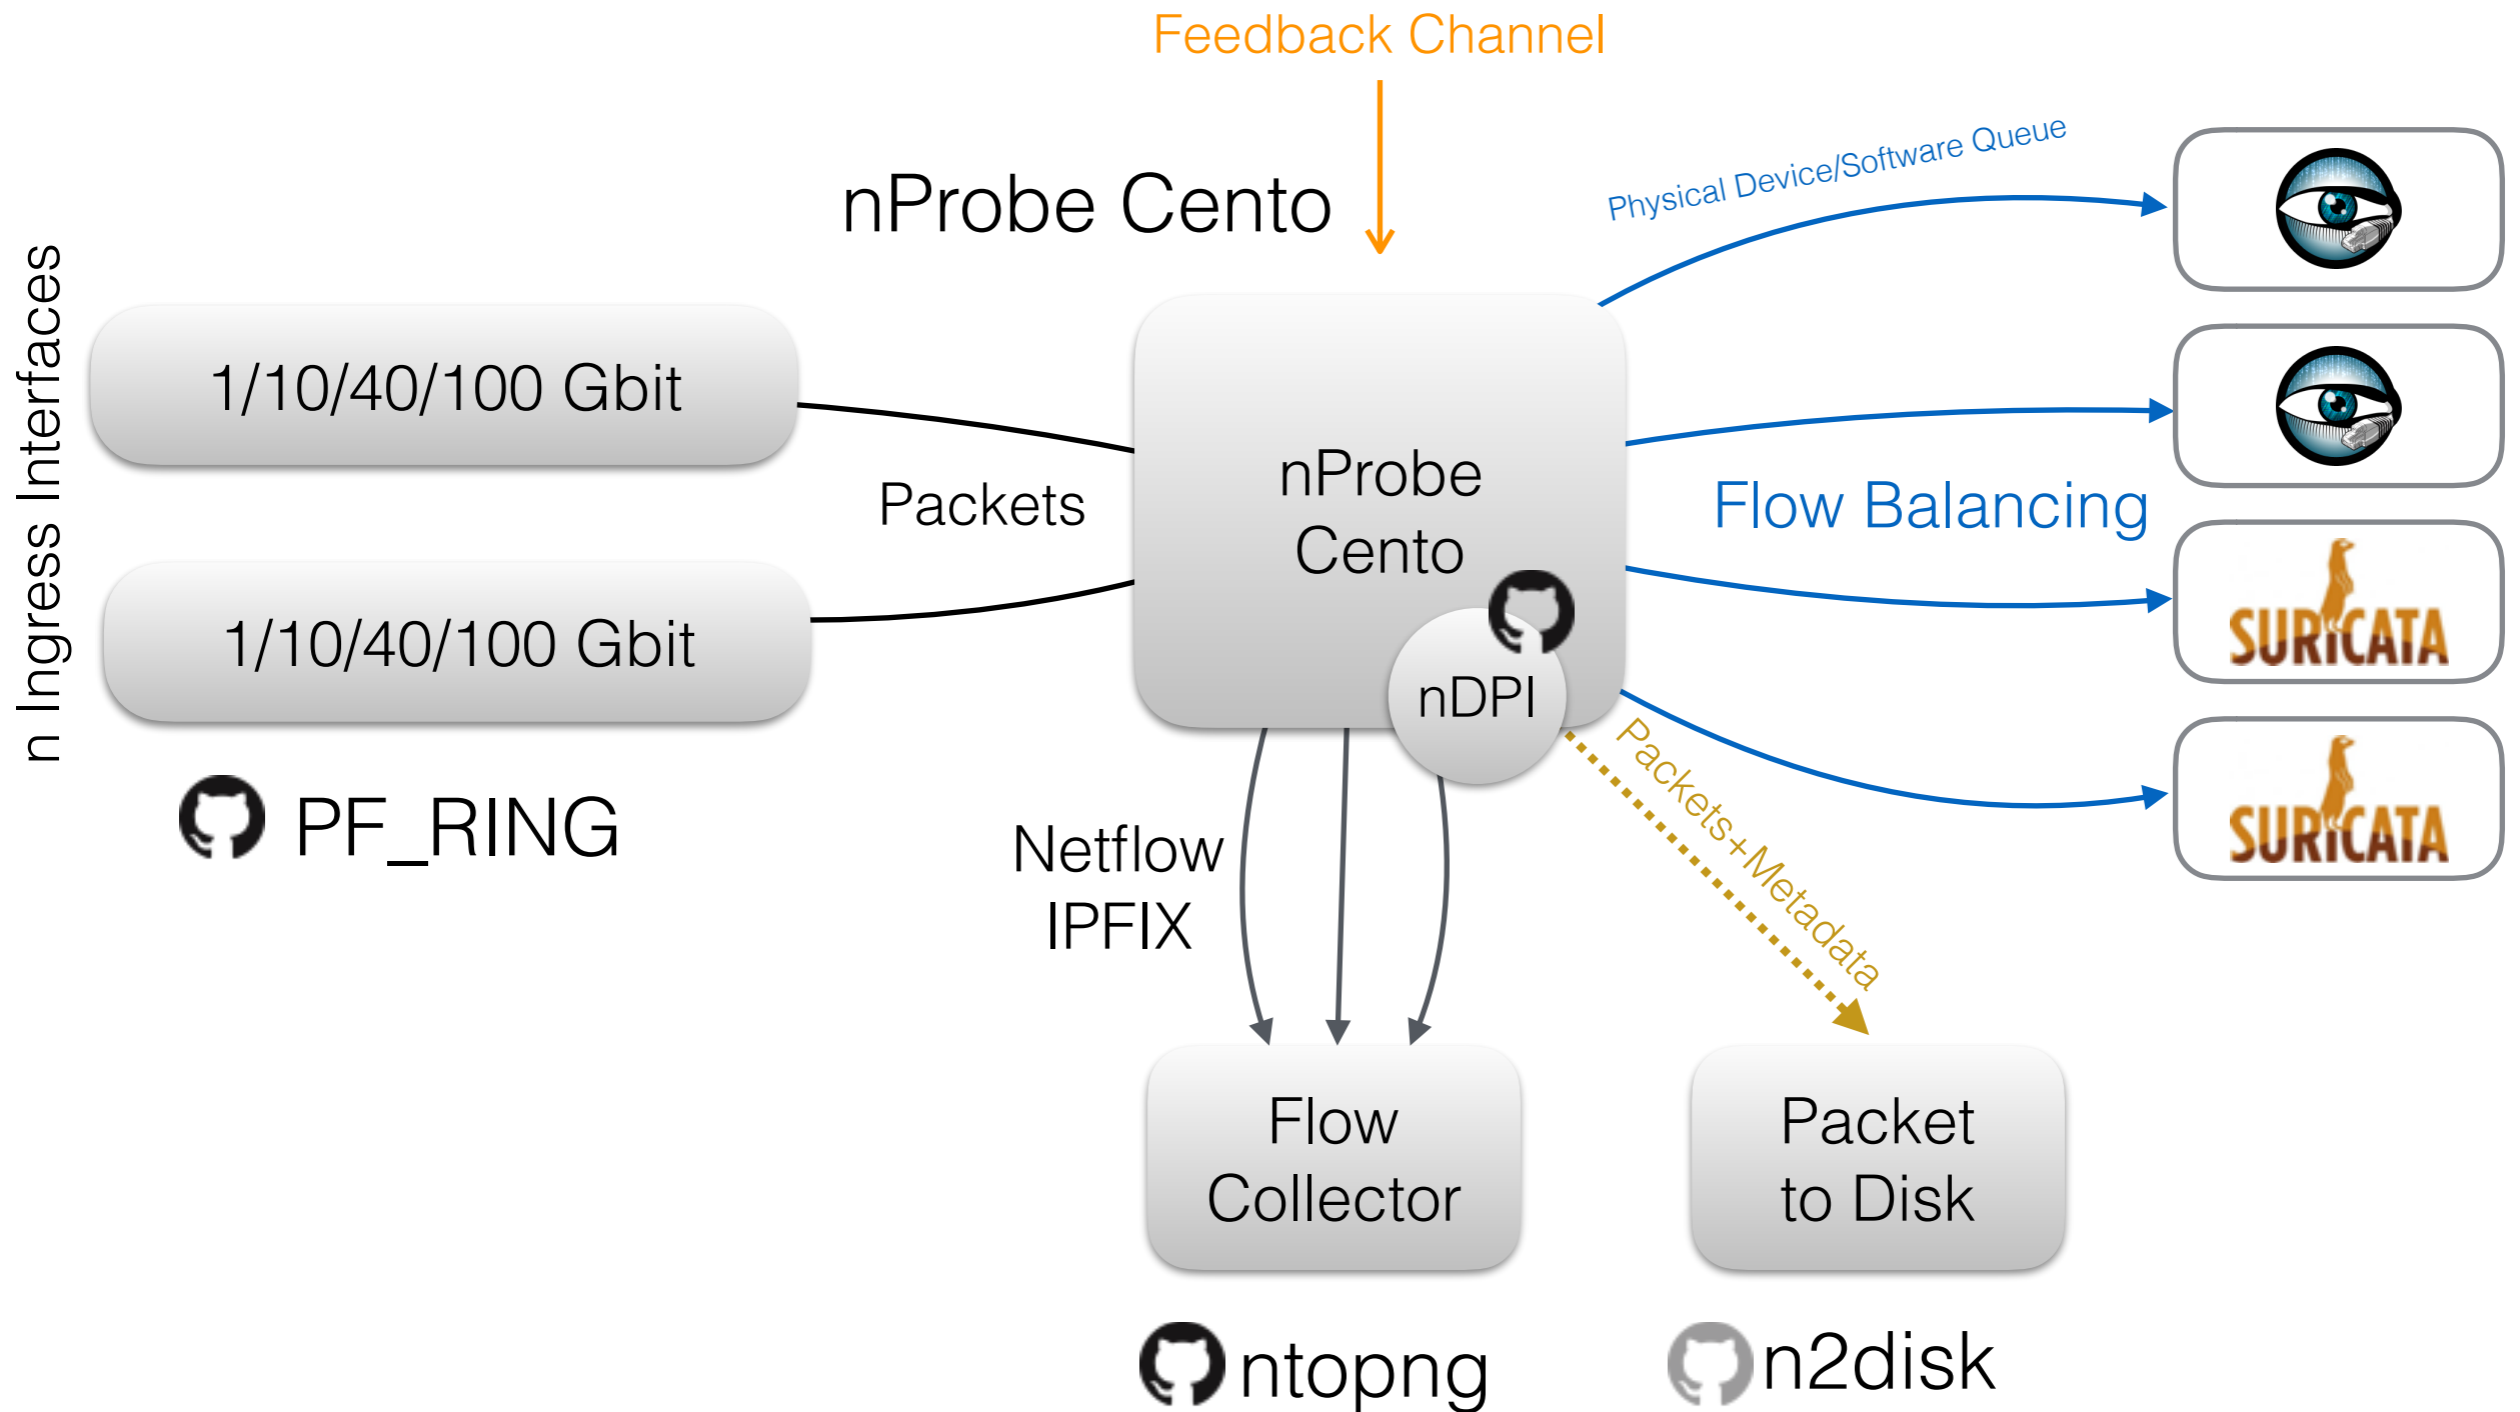
“Augmented” Flow Metadata [1/2]

- Augmented flows are delivered in IPFIX or JSON format and contain:
 - Standard flow fields (IPs, ports, bytes, packets, ...).
 - Client, server, and application latency.
 - Link/connection “quality” (packet OOO, retransmissions, fragmentation)
 - DNS/HTTP, and other application query data.
 - Application-specific timing (e.g. BitTorrent HashId)

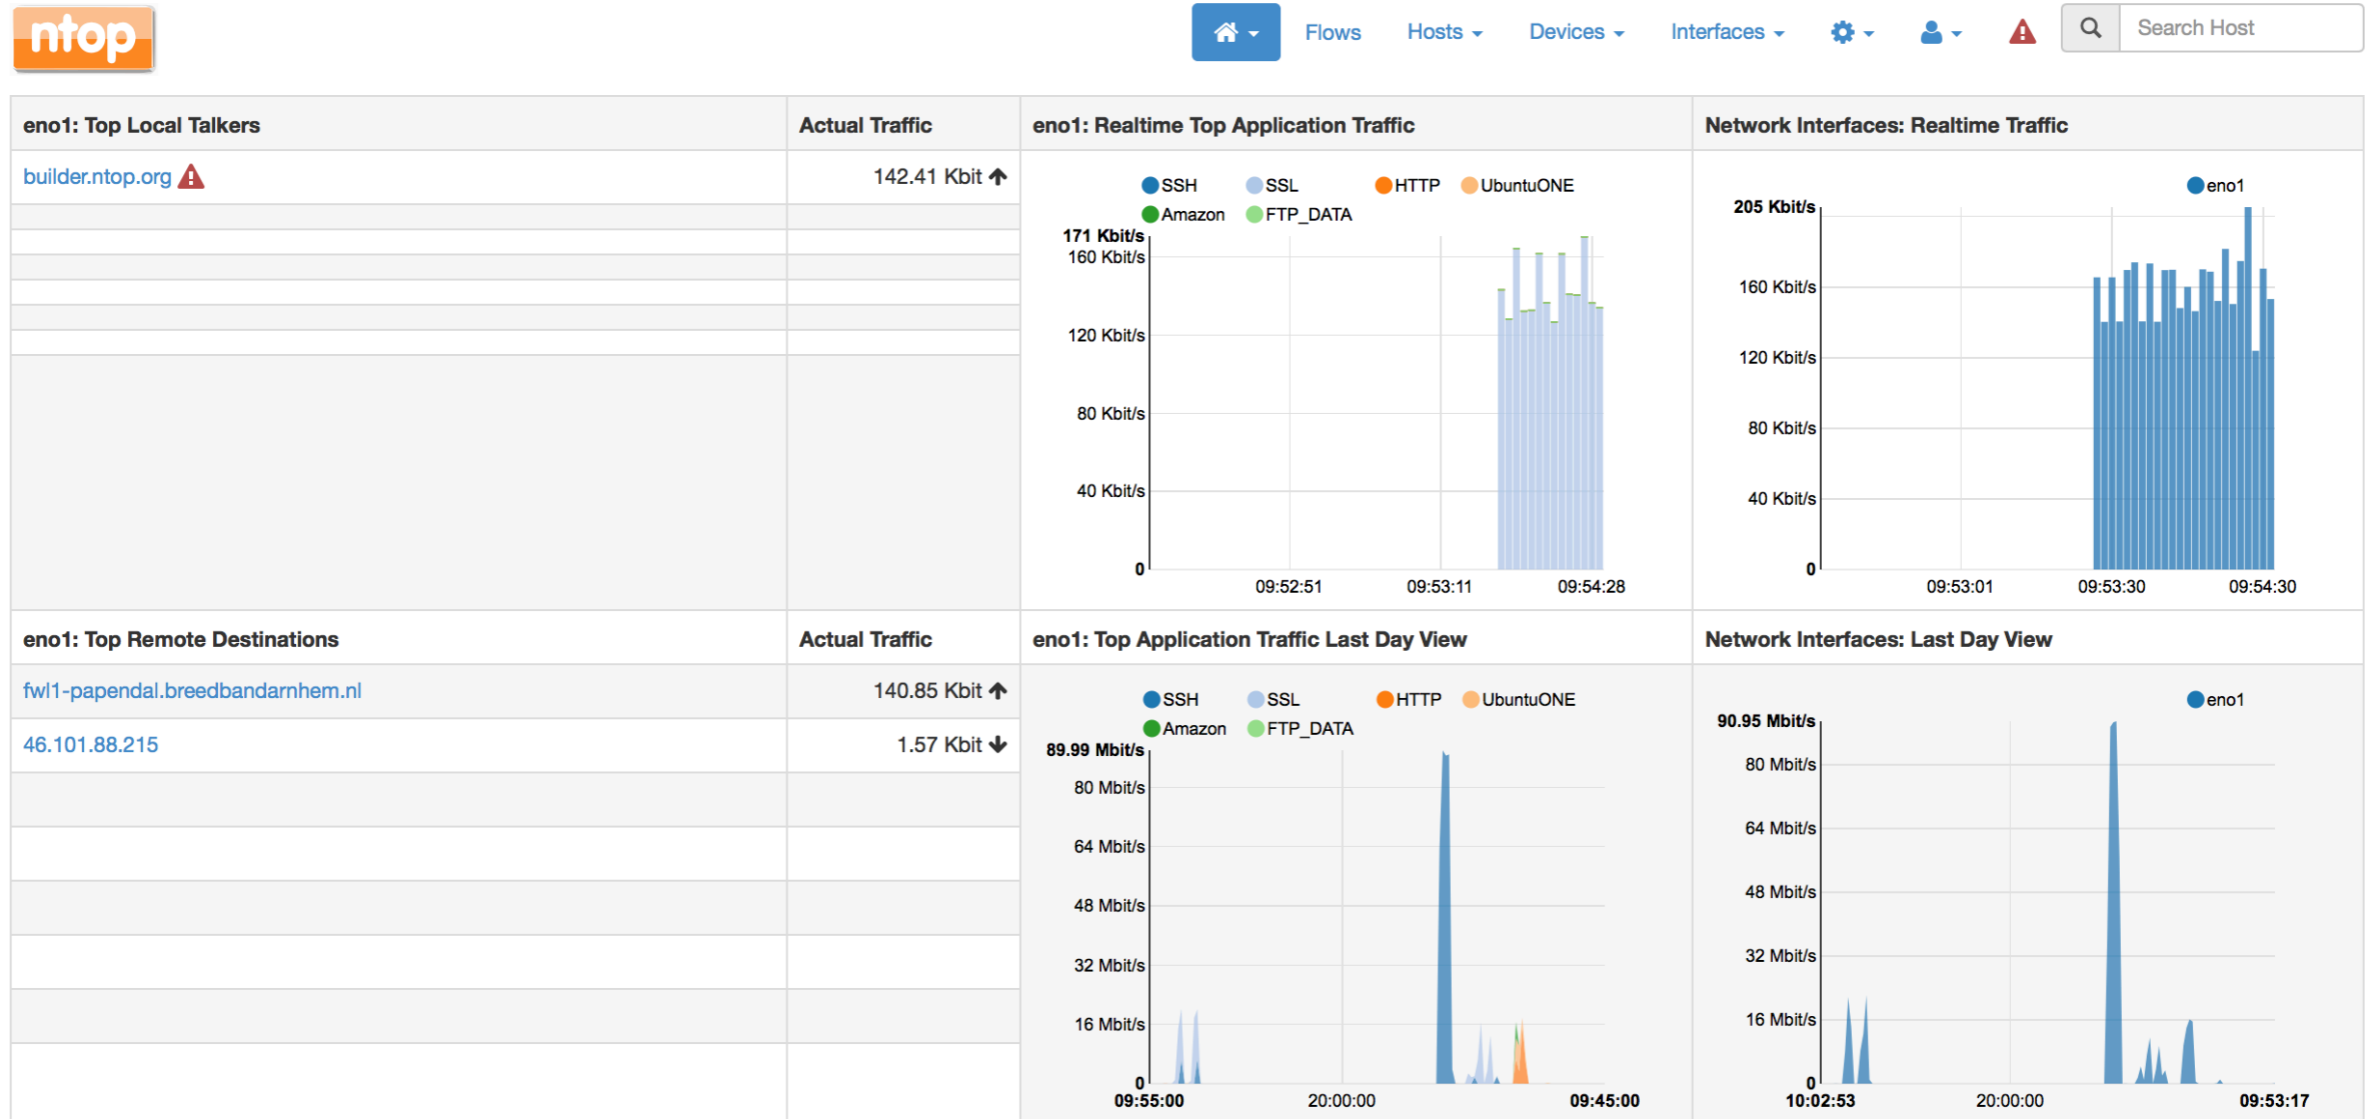
“Augmented” Flow Metadata [2/2]

- With this data, a good analytics platform can *detect and pinpoint*:
 - Network performance problems
 - Security attacks and break-ins
 - User-specific performance problems
- Don't forget that there are many activities (e.g. bots) that are hard to 'see' with just standard flow fields.

Putting It Together [1/2]



Putting It Together [2/2]



ntopng Enterprise v.2.5.161016
User [admin](#) Interface [eno1](#)

0% 42.09 Kbps [22 pps] 22.52 Kbps 15.34 Kbps

Uptime: 3 days, 20 h, 4 min, 21 sec
⚠ 6,609 Alerts 10 Hosts 2 Devices 84 Flows



Analytics [1/2]

Centralised analytics work for summaries and lower volume with ntopng

- Big data-friendly flow transport (JSON, Kafka)
- Larger-scale modern analytics platforms
- Real-time, un-aggregated (raw) flows
- Fuses flow, routing, and other data (hot for ISPs)
- Understand enriched/augmented flow
- Open to build and integrate with third-party applications

Analytics [2/2]

User Expectations for Enterprise Tools


- Instant data cubes
- 90 days (or more) of instantly accessible raw flows
- No aggregation, no roll-ups, no data cube fragility
- Ad-hoc, raw data query: 95%ile response < 2 seconds
- Interactive BGP path visualisations
- Instantaneous filtering and drill-down with no waiting

Analytics Example: Kentik Detect

**Secure Multi-Tenant SaaS
or
On-Premises Option**

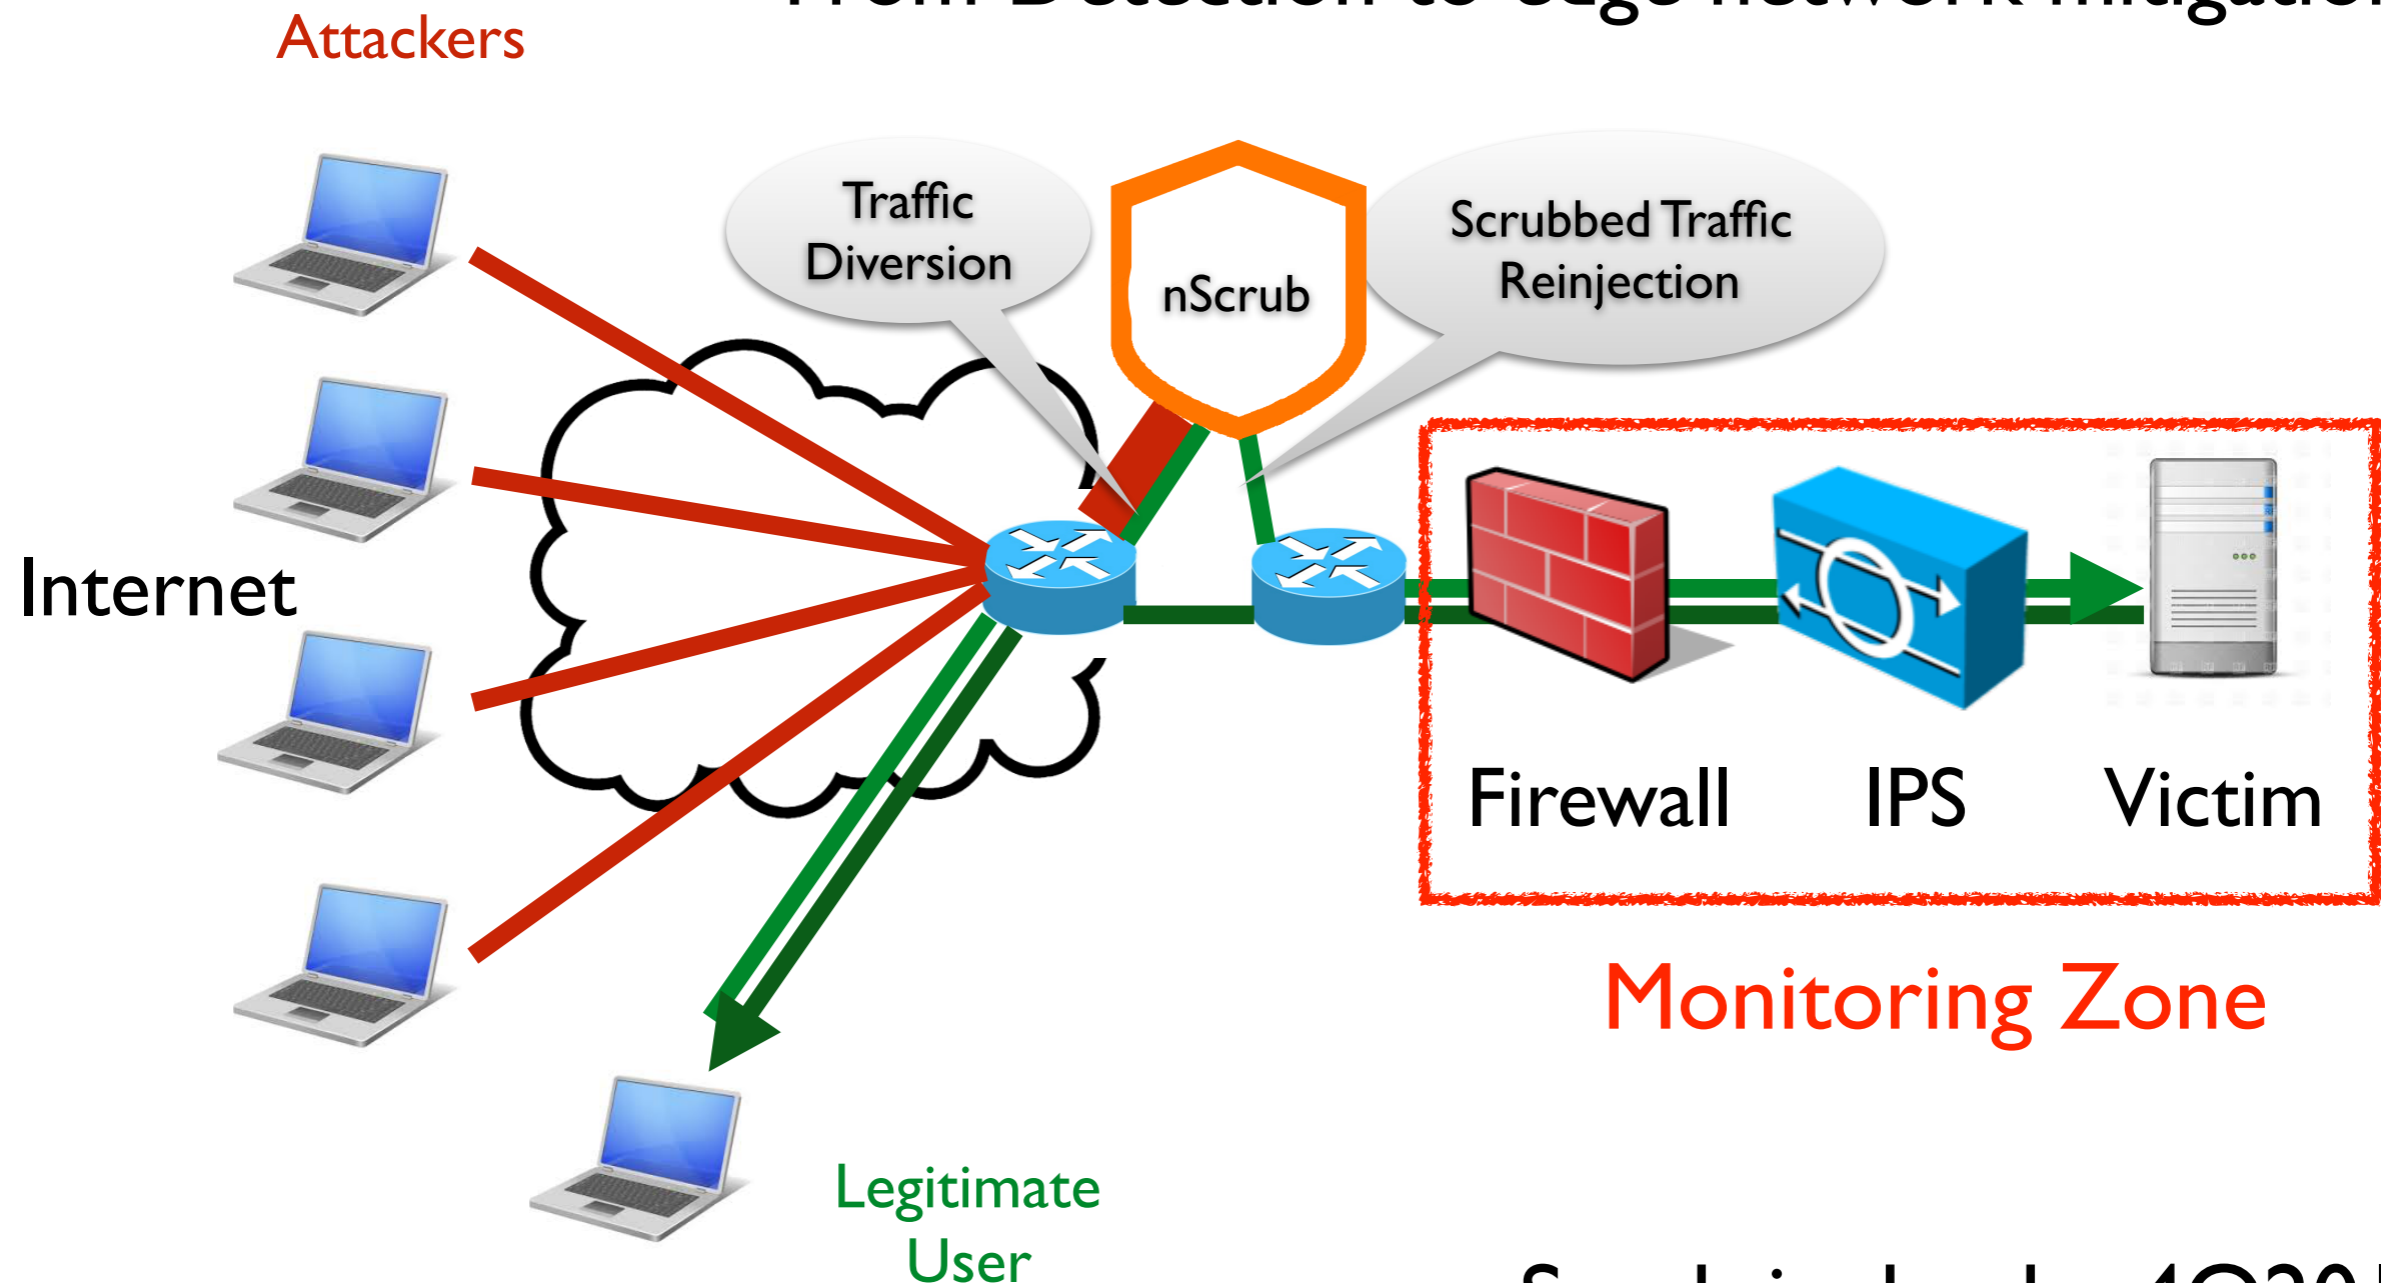
Network traffic &
performance data

Cross-Organizational
Network Visibility



DDoS Protection for Monitoring and Enterprises

From Detection to edge network mitigation...



nScrub is due by 4Q2016

Conclusions

- This talk presented solutions for network traffic monitoring able to satisfy needs of both small businesses, and large ISPs.
- Most tools are available on github. Commercial ones are free for education, research, no profit.
- Using commodity hardware and software optimised for modern computing architectures, it is possible to address many traffic monitoring issues with limited costs.
- We demonstrated that rich traffic monitoring at 10/40/100 Gbit is it now feasible (no more excuses for using SNMP interface traffic counters).