Affordable High-Speed Sensors Everywhere

ntop Meetup
Flocon 2016, Daytona Beach
Jan 13th 2016

ntop

FloCon 2016

Welcome to the ntop Meetup

- Meeting Goals:
  - Meet ntop users
  - Describe our ideas and plans for 2016
  - Hear your comments
  - Adapt the software roadmap based on the feedback received.
  - Exchange ideas.

Outlook

1. Introduction
2. 100 Gbit card support in PF_RING.
3. Flows for everything, pcap for something.
4. High-speed packet-to-disk for connected forensics.
5. Non-IPFIX flow export protocols.
6. Flow Monitoring and DDoS: software scrubbers for protecting networks and sensors.
7. Monitoring and Inline Traffic Policing.

Affordable High-Speed Sensors Everywhere

- Network monitoring is often perceived as a costly activity as it requires network sensors to be deployed where network traffic flows.
- NetFlow/IPFIX enable people to collect network metrics in an open format (sometimes not fully following the standard, e.g. Cisco ASA).
- Standard evolve very slowly, vendors often report just basic metrics, people need new features (e.g. DPI), then custom sensors need to be deployed.
- Requirement: sensors must be affordable and rich in measurement metrics (we need more than bytes/packets).

Affordable **High-Speed** Sensors Everywhere

- What is high-speed today?
  - 1 Gbit for home and small offices
  - 10 Gbit for medium business
  - 100 Gbit for ISP/large business
- Often lines are not fully occupied, but the number of flows/sec increase as many companies rely on cloud services that replaced LAN-based services.
- In addition to speed, sensors need to provide rich application-based metrics as more speed does not mean poor metrics.

Affordable High-Speed Sensors **Everywhere**

- In order to make network monitoring commodity, all the traffic must be monitored, not just the core network.
- What is a price ballpark for a line-rate sensor (hardware + software) that could be placed everywhere on a network?
  - 1 Gbit:1'000 USD
  - 2 x10 Gbit: 2'500 USD
  - 100 Gbit: 25'000 USD
- This is our goal for 2016.

100 Gbit Support in PF_RING

- As PF_RING is becoming increasingly popular (tools like YAF, Bro, Snort, Suricata support it just to mention a few), we have decided to add as many network adapters as possible so that you can code your app once, and deploy it on top of various NICs, just changing the device name.
- Currently PF_RING supports the following adapters:

Flows for everything, pcap for something [1/3]

- Most organizations need to have evidence of all activities happened in their network.
- Many companies 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 [2/3]

- 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.
- In essence recording all traffic to disk is not a good idea. Instead recoding all the interesting traffic to disk is a good idea.

Flows for everything, pcap for something [3/3]

- How we plan to do it:
  - Generate augmented flows (e.g. include HTTPS host name, or DNS query) on all traffic, so that there is evidence of all activities happened on the network.
  - Record traffic more efficiently:
    - Discard/slice un-necessary protocols.
    - Full packets for some hosts (e.g. core servers or hosts that produce security alerts).

Towards a More Efficient DPI [1/2]

- Since 4 years, ntop develops an open-source DPI library named nDPI (more than 200 protocols supported).
- While nDPI is pretty efficient (with 2 cores it can handle 10 Gbit of traffic), “flows for everything” do not need advanced DPI, but very efficient DPI as all the traffic has to be turned into flows.
- As most “modern” Internet traffic is HTTP-based, we have decided it was time to develop a faster yet more limited DPI toolkit. This is why we developed µ-nDPI.

Towards a More Efficient DPI [2/2]

- µ-nDPI main goals are:
  - Efficiency: the performance degradation and memory overhead should be almost unnoticeable.
  - Focused: (initially?) handle only HTTP/SSL/DNS and subprotocols (e.g. HTTP.Facebook).
  - Metadata: provide limited metadata (e.g. DNS Query and reply code) enough to identify issues that need to be further analysed.
- “pcap for something” will leverage on µ-nDPI.

High-speed Packet-to-Disk for Connected Forensics [1/2]

- n2disk is ntop's 2x10Gbit packet-to-disk application.
- Currently it is an independent component devoted to write traffic to disk. However its integration with flow-monitoring tools is necessary.
- We plan to:
  - Write flowId+appId into the n2disk packet index.
  - Attach n2disk as a consumer of a "first stage" flow-generation application.
  - Avoid writing unnecessary packets.

High-speed Packet-to-Disk for Connected Forensics [2/2]

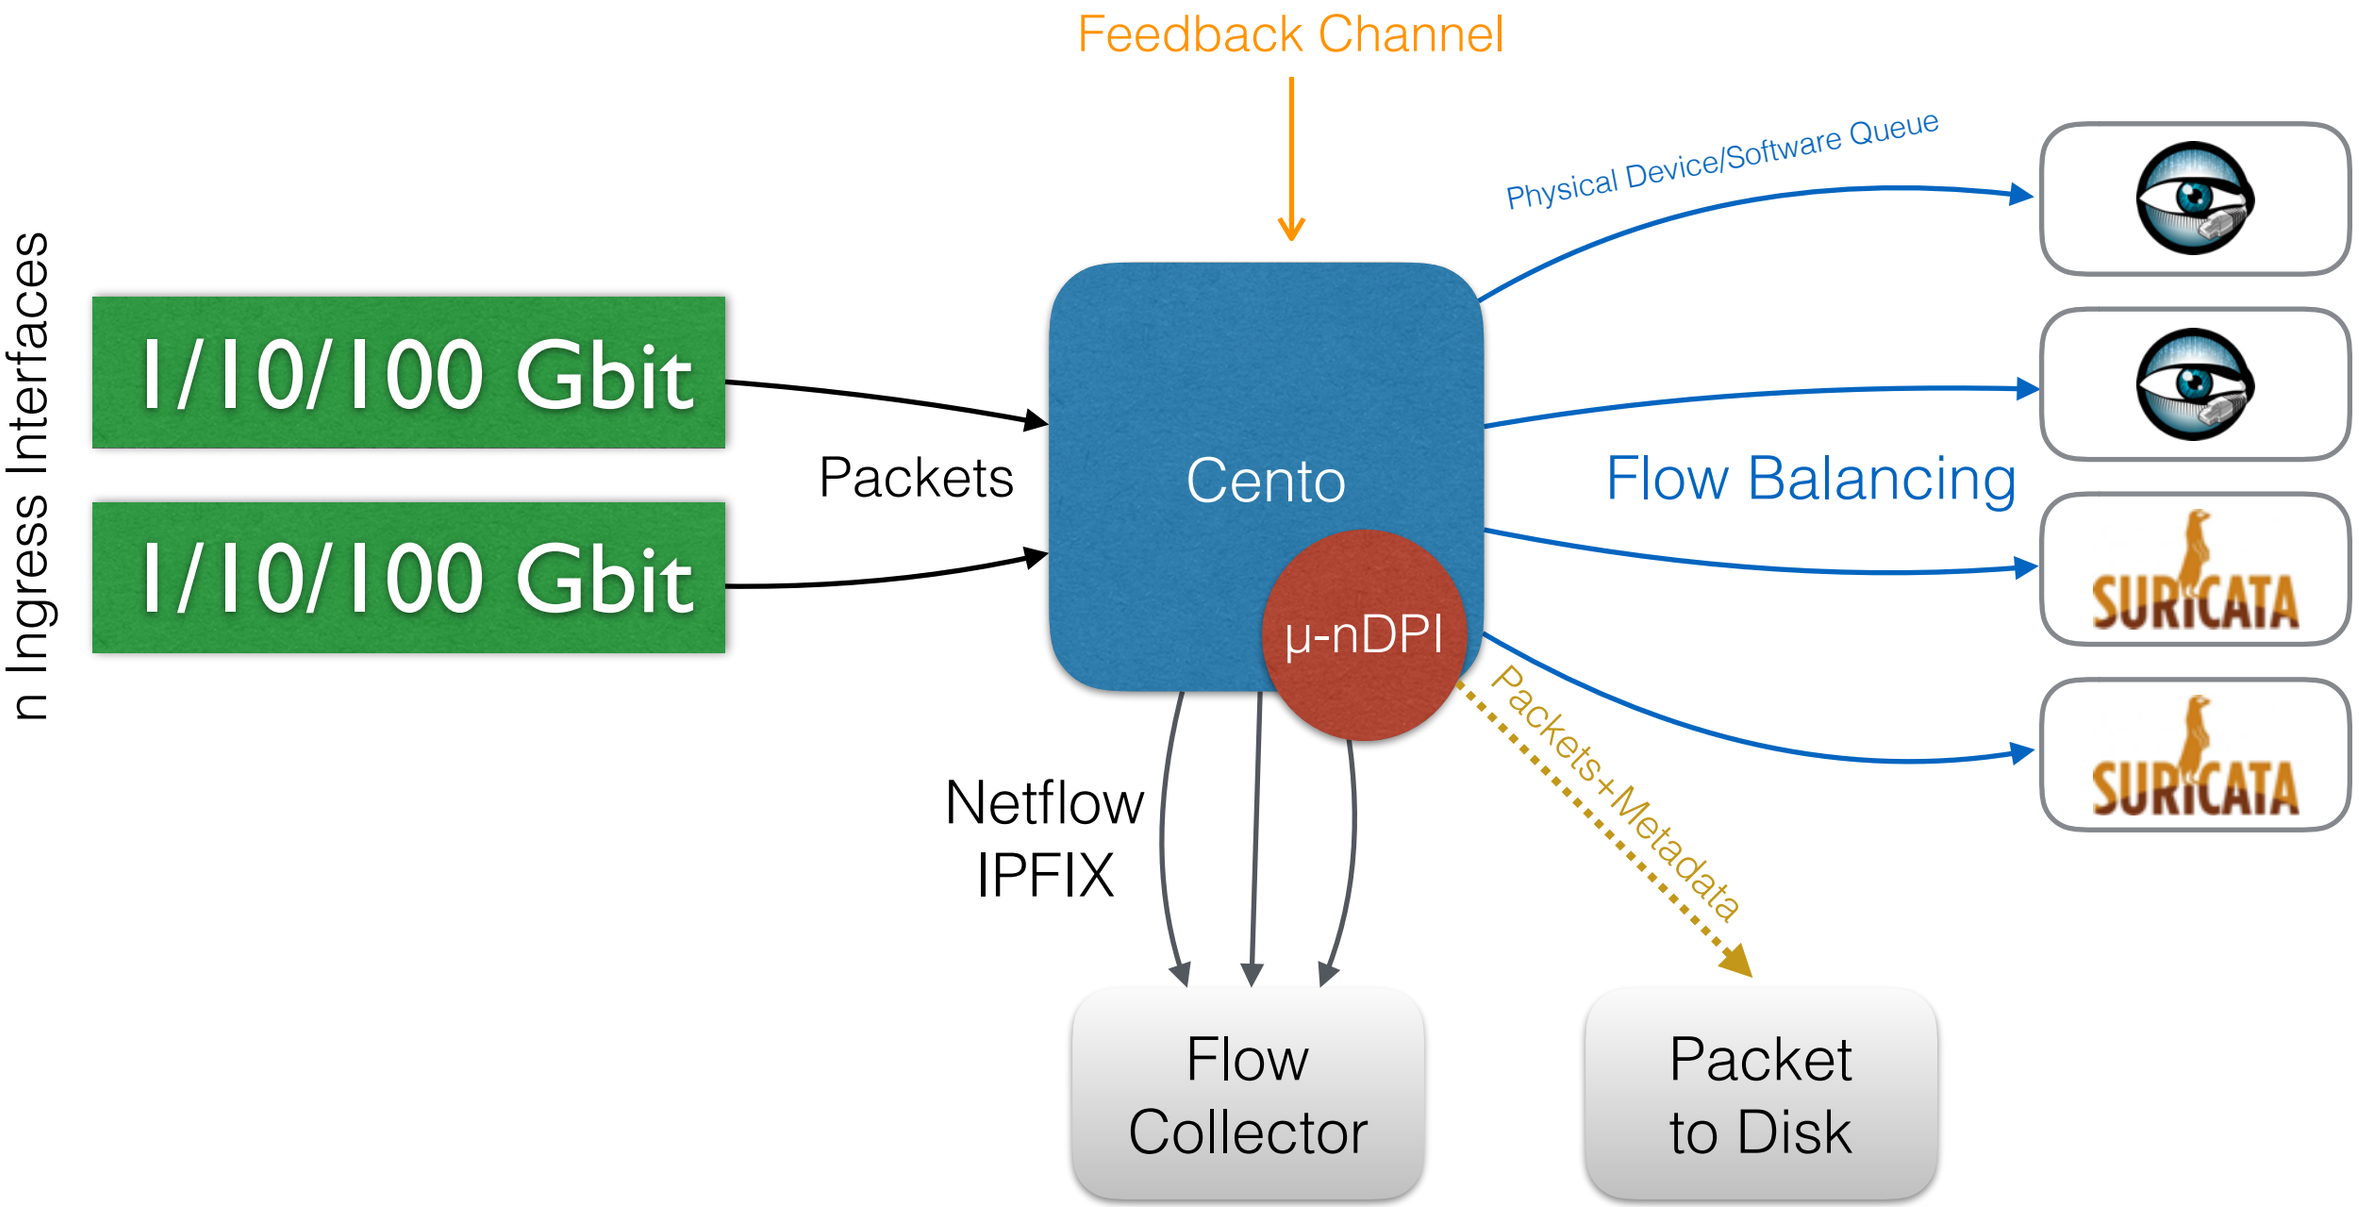

Non-IPFIX Flow Export Protocols [1/2]

- While the flow paradigm is consolidated, NetFlow/IPFIX formats have several limitations (e.g. missing compression, issues with variable-length data, complex template-based format).
- Many big-data systems can handle JSON or binary format (e.g. ProtocolBuffers) but not NetFlow/IPFIX.
- The use of Apache Kafka is also becoming widespread and a non-NetFlow/IPFIX format is becoming important.

Non-IPFIX Flow Export Protocols [2/2]

- All ntop apps currently support JSON as export format (open format, but not too efficient) but it is necessary to also add a more efficient/compressed binary format for flow exchange.
- Alternatives:
  - Kentik has proposed kflow a binary format based on Cap'n Proto (efficient successor of ProtocolBuffers), that is worth to consider as addition.
  - MessagePack (binary JSON) can also be considered but contrary to kflow it does not have a fixed flow format.

Flow Monitoring and DDoS [1/2]

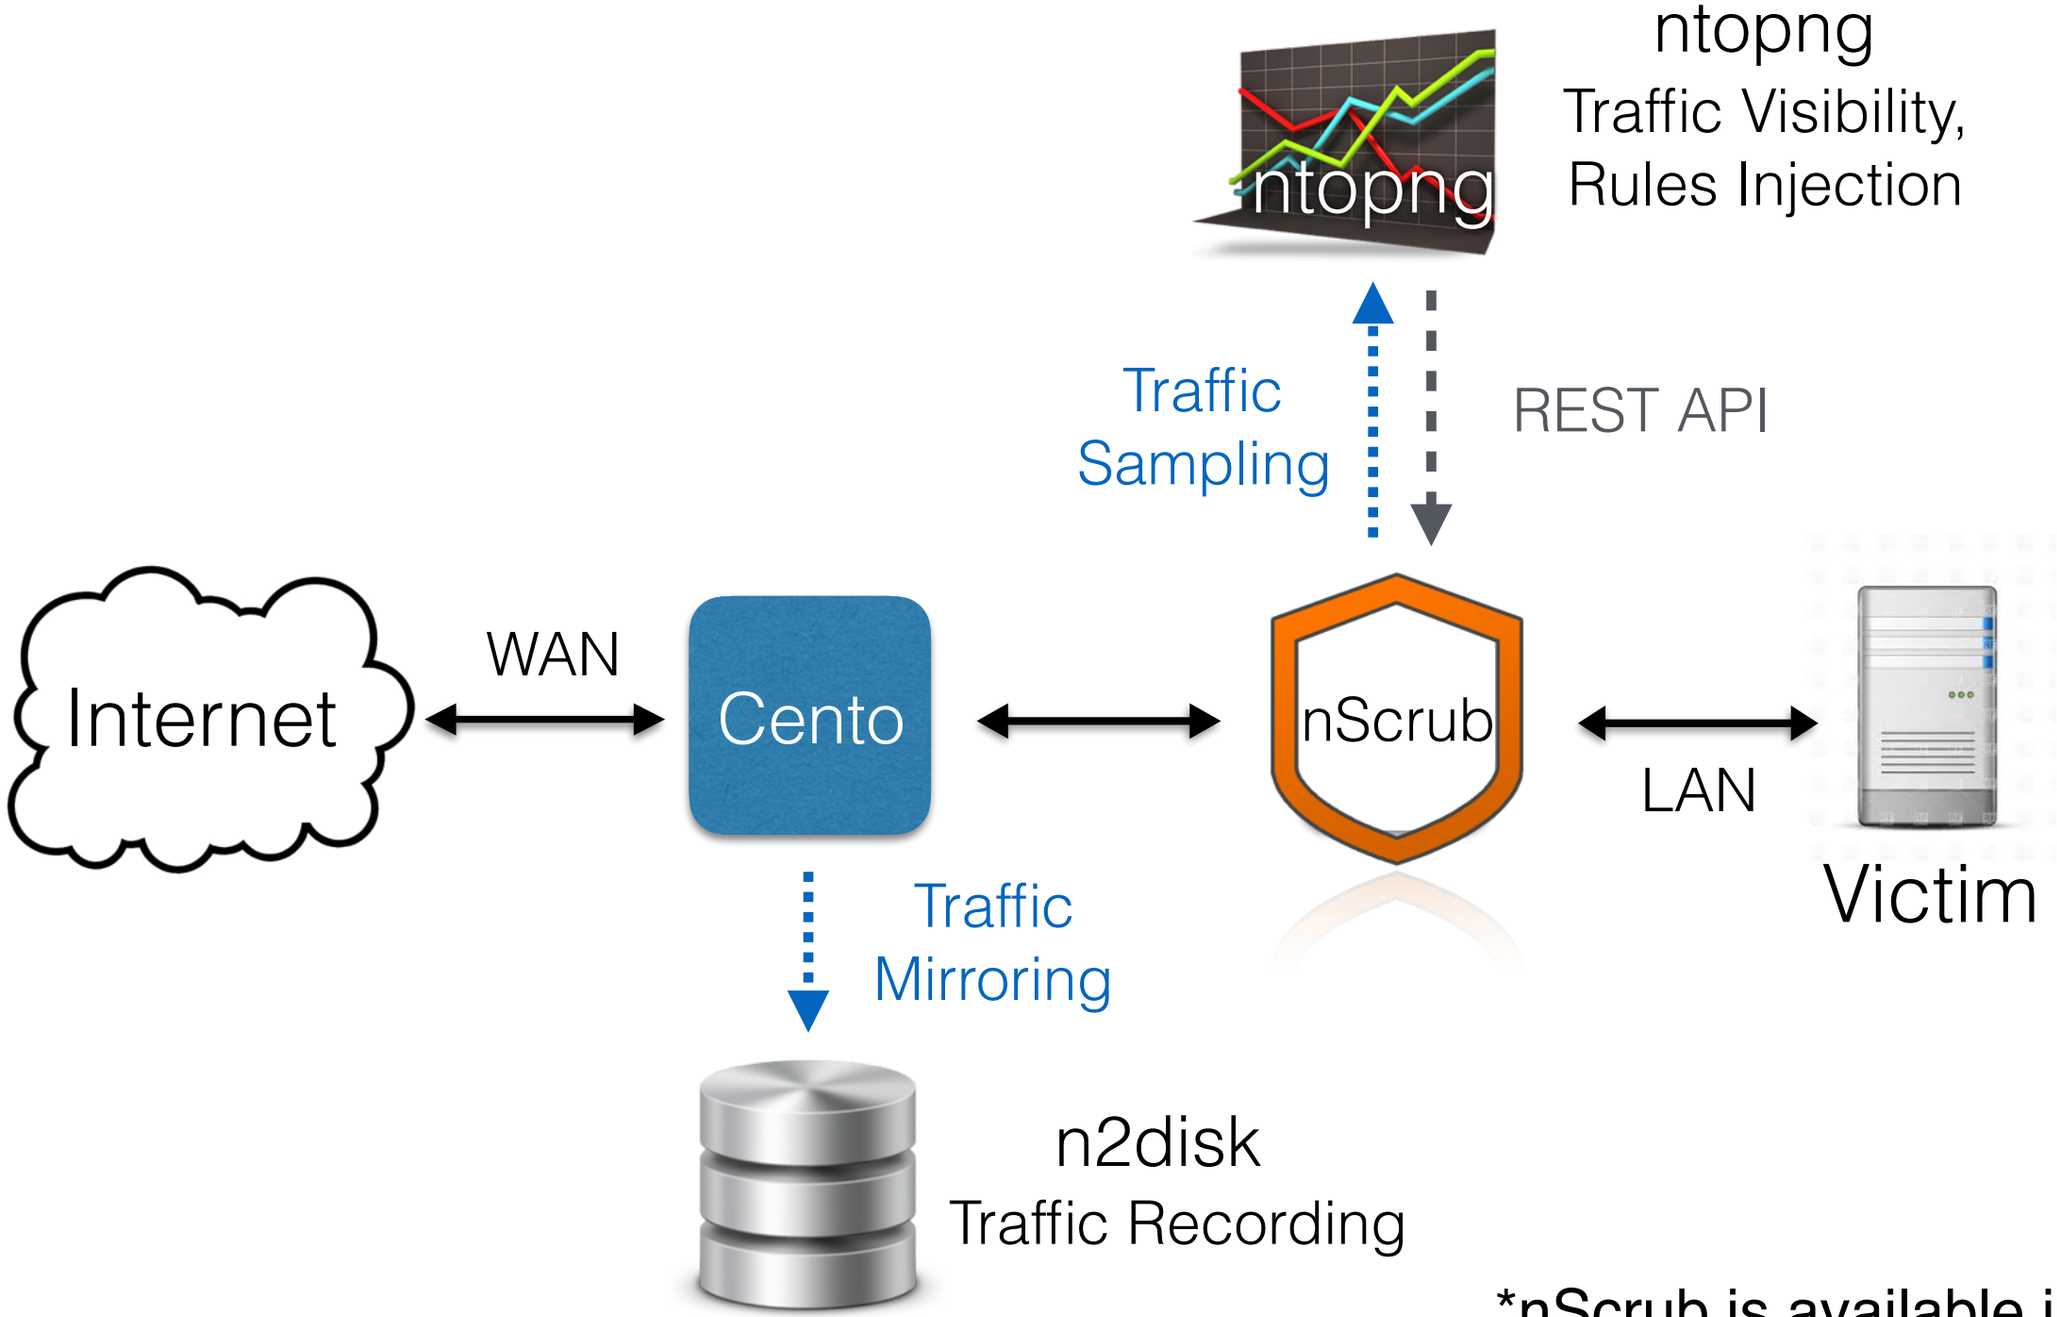

*nScrub is available in beta

Flow Monitoring and DDoS [2/2]

Flow Monitoring Mode

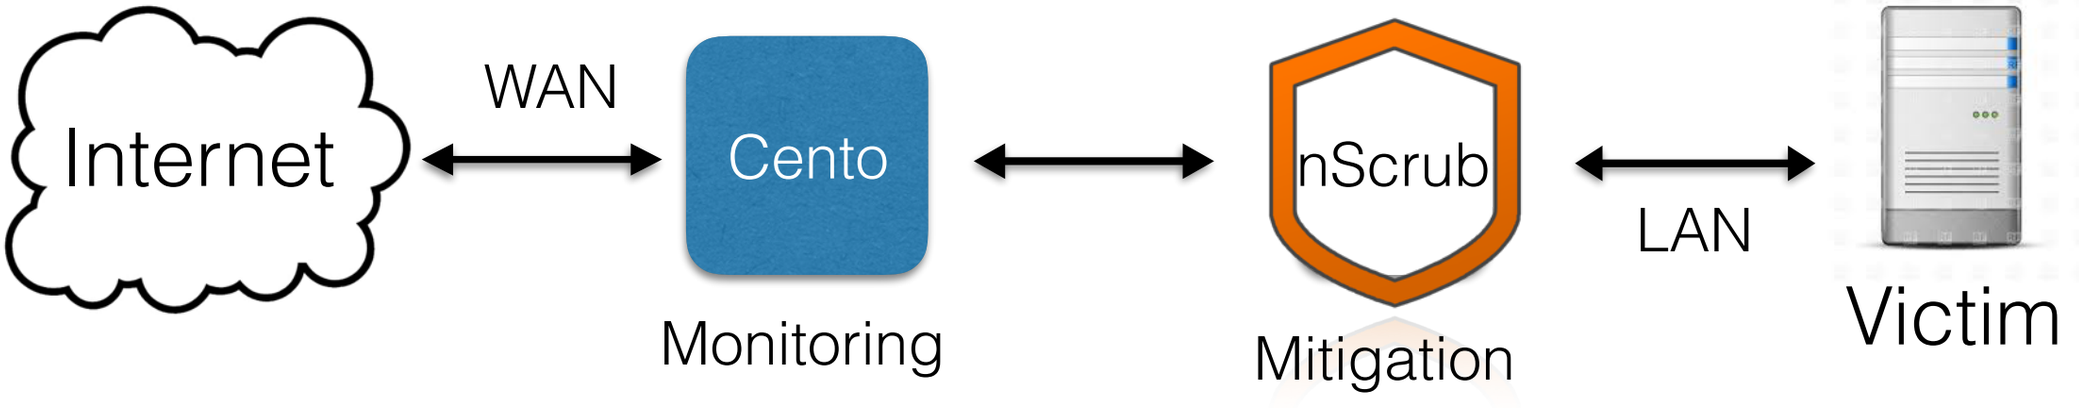

Layer-7 Filtering Mode

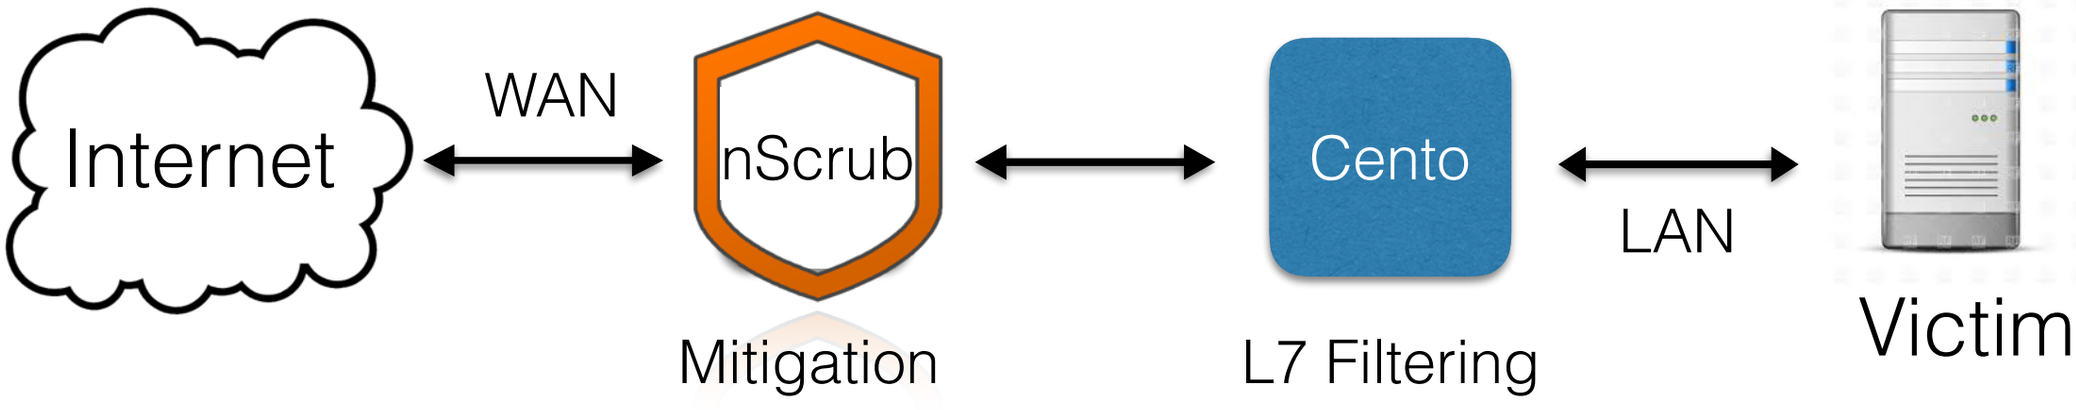

Monitoring and Inline Traffic Policing [1/2]

- In small/medium enterprises the use of independent components (e.g. probe, collector, packet-to-disk) creates operational issues.
- Companies prefer to collapse functionalities previously implemented by several tools into a single application.
- Examples include drop of:
  - Application protocols (e.g. no Skype, BitTorrent).
  - Selected web-sites URLs (e.g. porn or compromised sites).

Monitoring and Inline Traffic Policing [2/2]

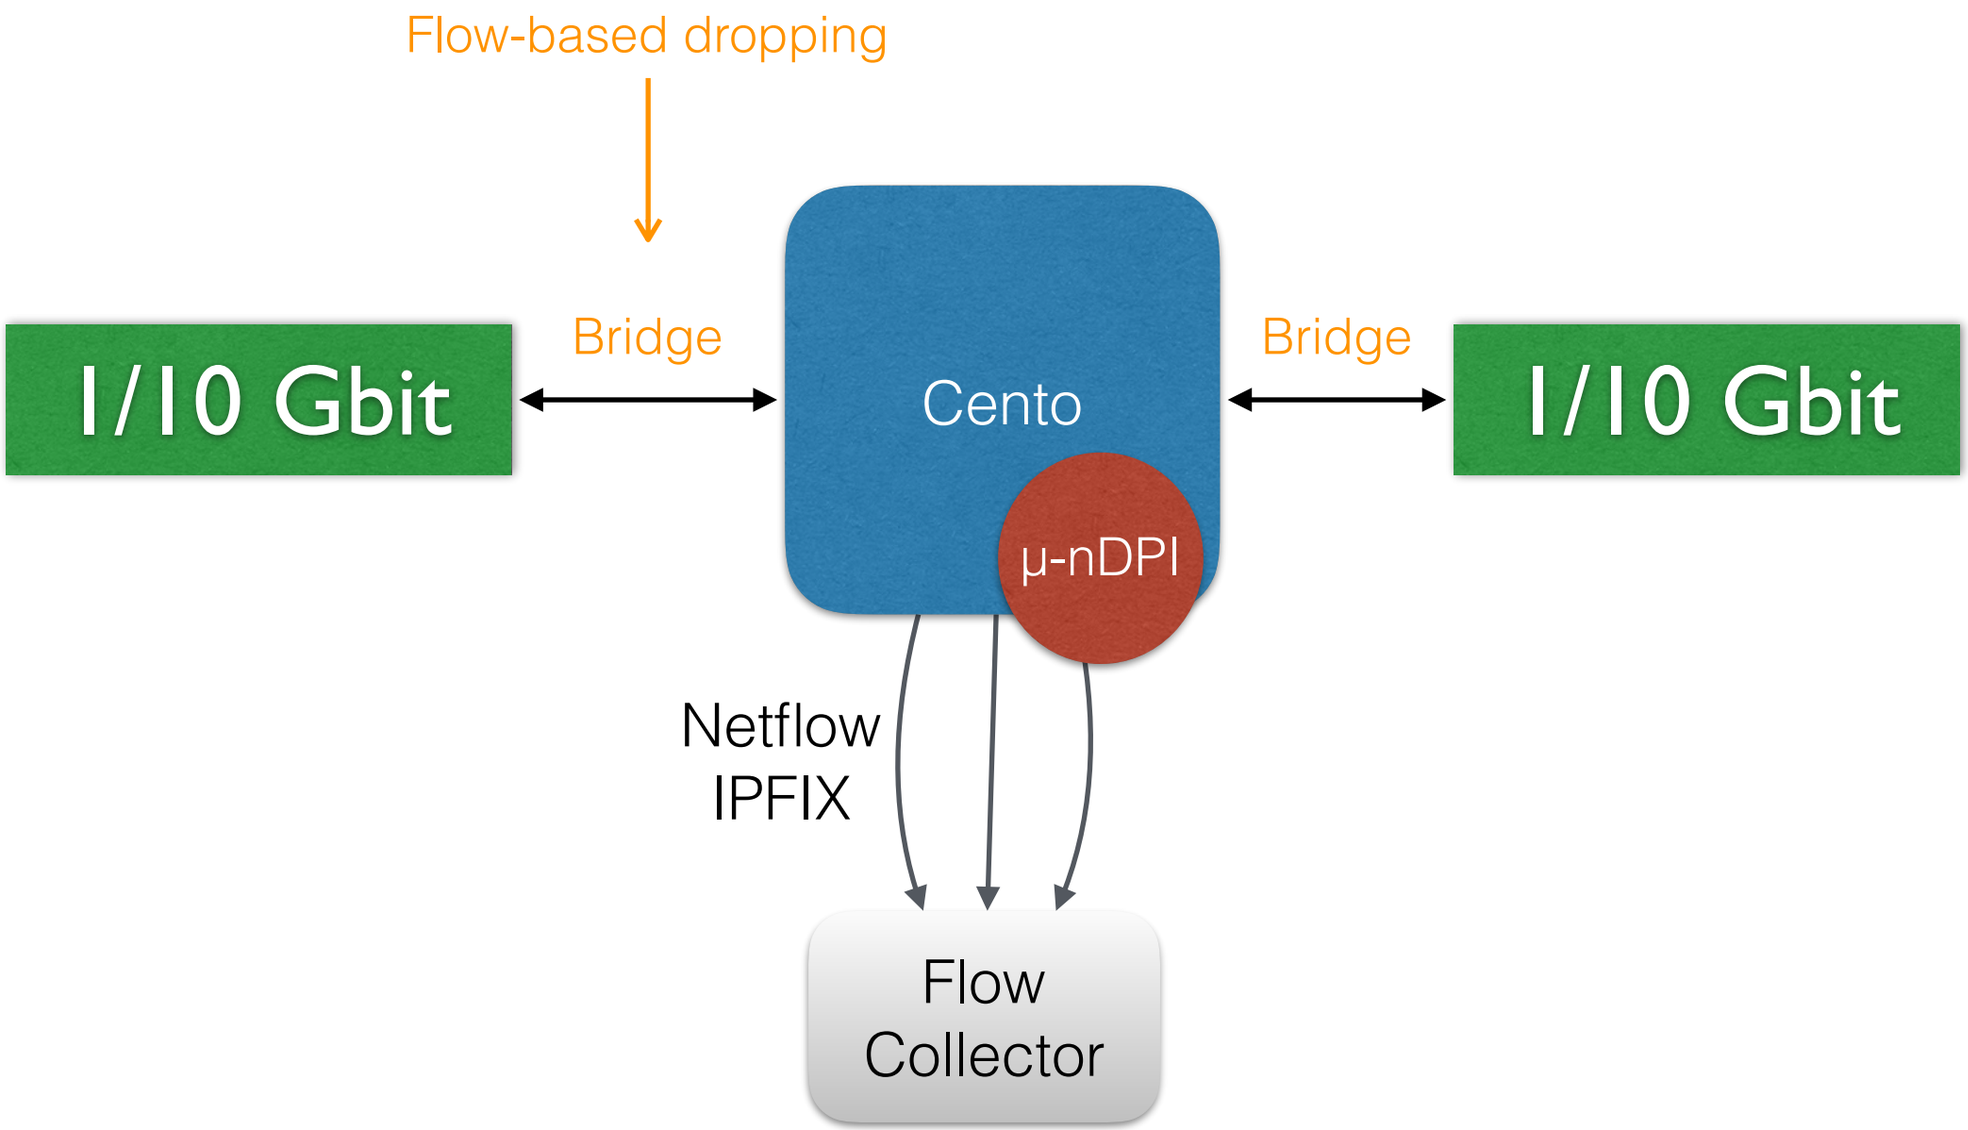

Some Open Questions: Hardware

- What features PF_RING should support in order to better exploit hardware adapters?
- Can flow pre-computation/generation be (partially?) offloaded to a FPGA-based network adapter?
- IDS(-like) applications:
  - Would an IDS exploit packet parsing done in PF_RING (or onto a NIC) for accelerating processing?
  - Shall PF_RING implement a stateful flow-processor for discarding selected flows before they hit the IDS (similar to snort DAQ but implemented in PF_RING)?

Some Open Questions: Software

- As encrypted traffic increases, is DPI still a good idea ? If not, what shall we do?
- As network speed increases (and time to process packets decreases), what is the minimum set of metrics a high-speed sensor shall compute?
- In order to avoid un-necessary processing, what an IDS can expect from a flow-preprocessor in terms of features?
- Is there any (open) feedback mechanism for IDSs to report apps (e.g. cento or nScrub) about detected threats so to that packets could be dropped?

Some Open Questions: Analytics

- Should we build tools that analyse and extracts actionable insights from the great deal of data we are producing?
- E.g., Network Behaviour Analysis (NBA) / Network Behaviour Anomaly Detection (NBAD)
- What kind of insights are desirable?
  - Policy violations
  - Intrusion detection
  - …

Use Cases and Roadmap Discussion

- Kentik
- Suricata Team
- Accolade Technology
- PacketChaser/FireEye
- Principal
- Napatech
- Telesoft
- Amazon
- ntop Users