# Flow-Based Network Monitoring using nProbe and ntopng

Simone Mainardi, PhD
@simonemainardi
mainardi@ntop.org

# Agenda

- About ntop
- Flow-based network monitoring, beyond SNMP
- nProbe: NetFlow/IPFIX/sFlow probe and collector
- ntopng: Web-Based visualization

# About ntop

- Private company devoted to development of Open Source network traffic monitoring applications.
- R&D Italy, Sales Switzerland.
- ntop (circa 1998) is the first app we released and it is a web-based network monitoring application.

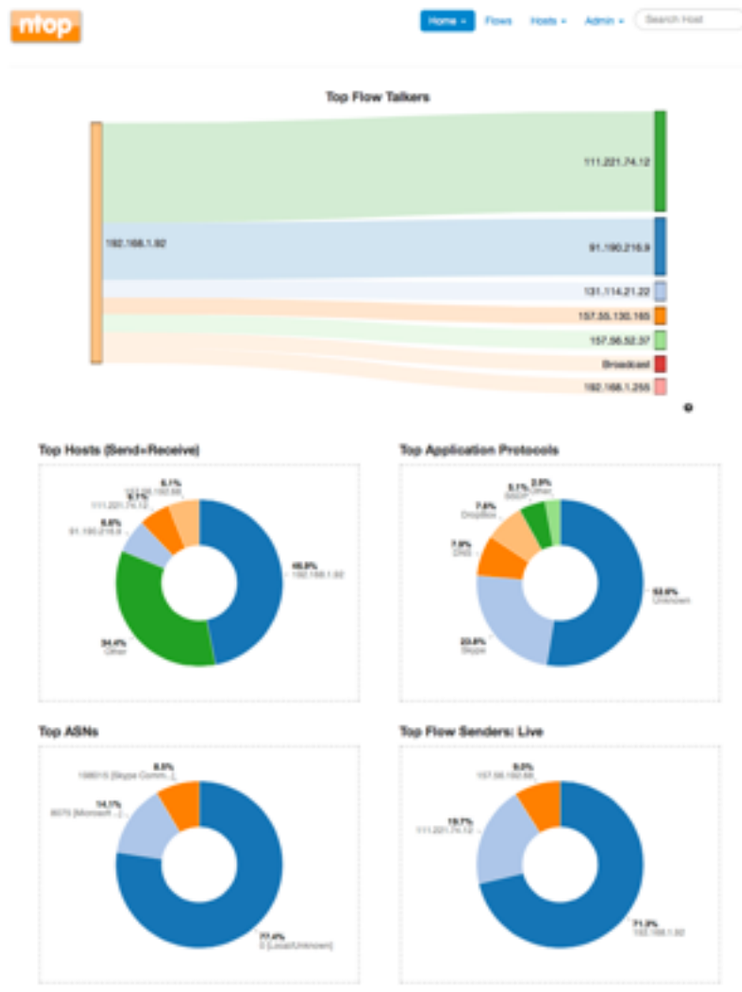

# Some Products we Developed [1/2]

- Our software is powering many commercial products...

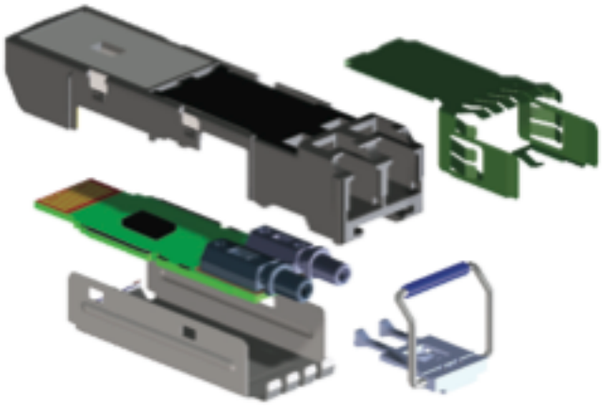

Integrated ASIC with JDSU technology


# Some Products we Developed [2/2]

- ...and allows packets to be received and transmitted at 1/10 Gbit line rate with no loss, any packet size on Intel-based commodity NICs.
- So we accelerate not just our applications but also third party open source solutions including:



# Product Lines

- Open Source (https://github.com/ntop)
  - ntopng: Web-based monitoring application
  - PF_RING: Accelerated RX/TX on Linux
  - nDPI: Deep Packet Inspection Toolkit
- Proprietary
  - PF_RING ZC: 1/10/40/100 Gbit Line rate.
  - nProbe: 10G NetFlow/IPFIX Probe
  - nProbe Cento: flows+packets+security
  - n2disk/disk2n Network-to-disk and disk-to-network.
  - nScrub: Software DDoS Mitigation

# Motivation

- Without measurements we cannot evaluate the results, introduce improvements, quantify the success of our policies.

"If you can't measure it,  you can't improve it"
(Lord Kelvin, 1824 – 1907)

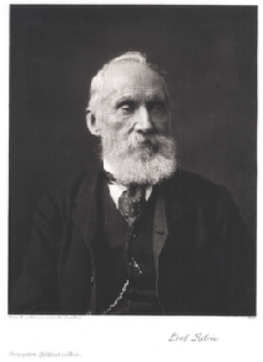

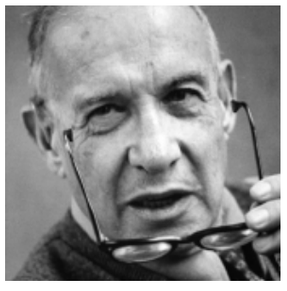

"If you can't measure it,  you can't manage it"
(Peter Drucker, 1909 – 2005)

# Network Monitoring with Check_MK

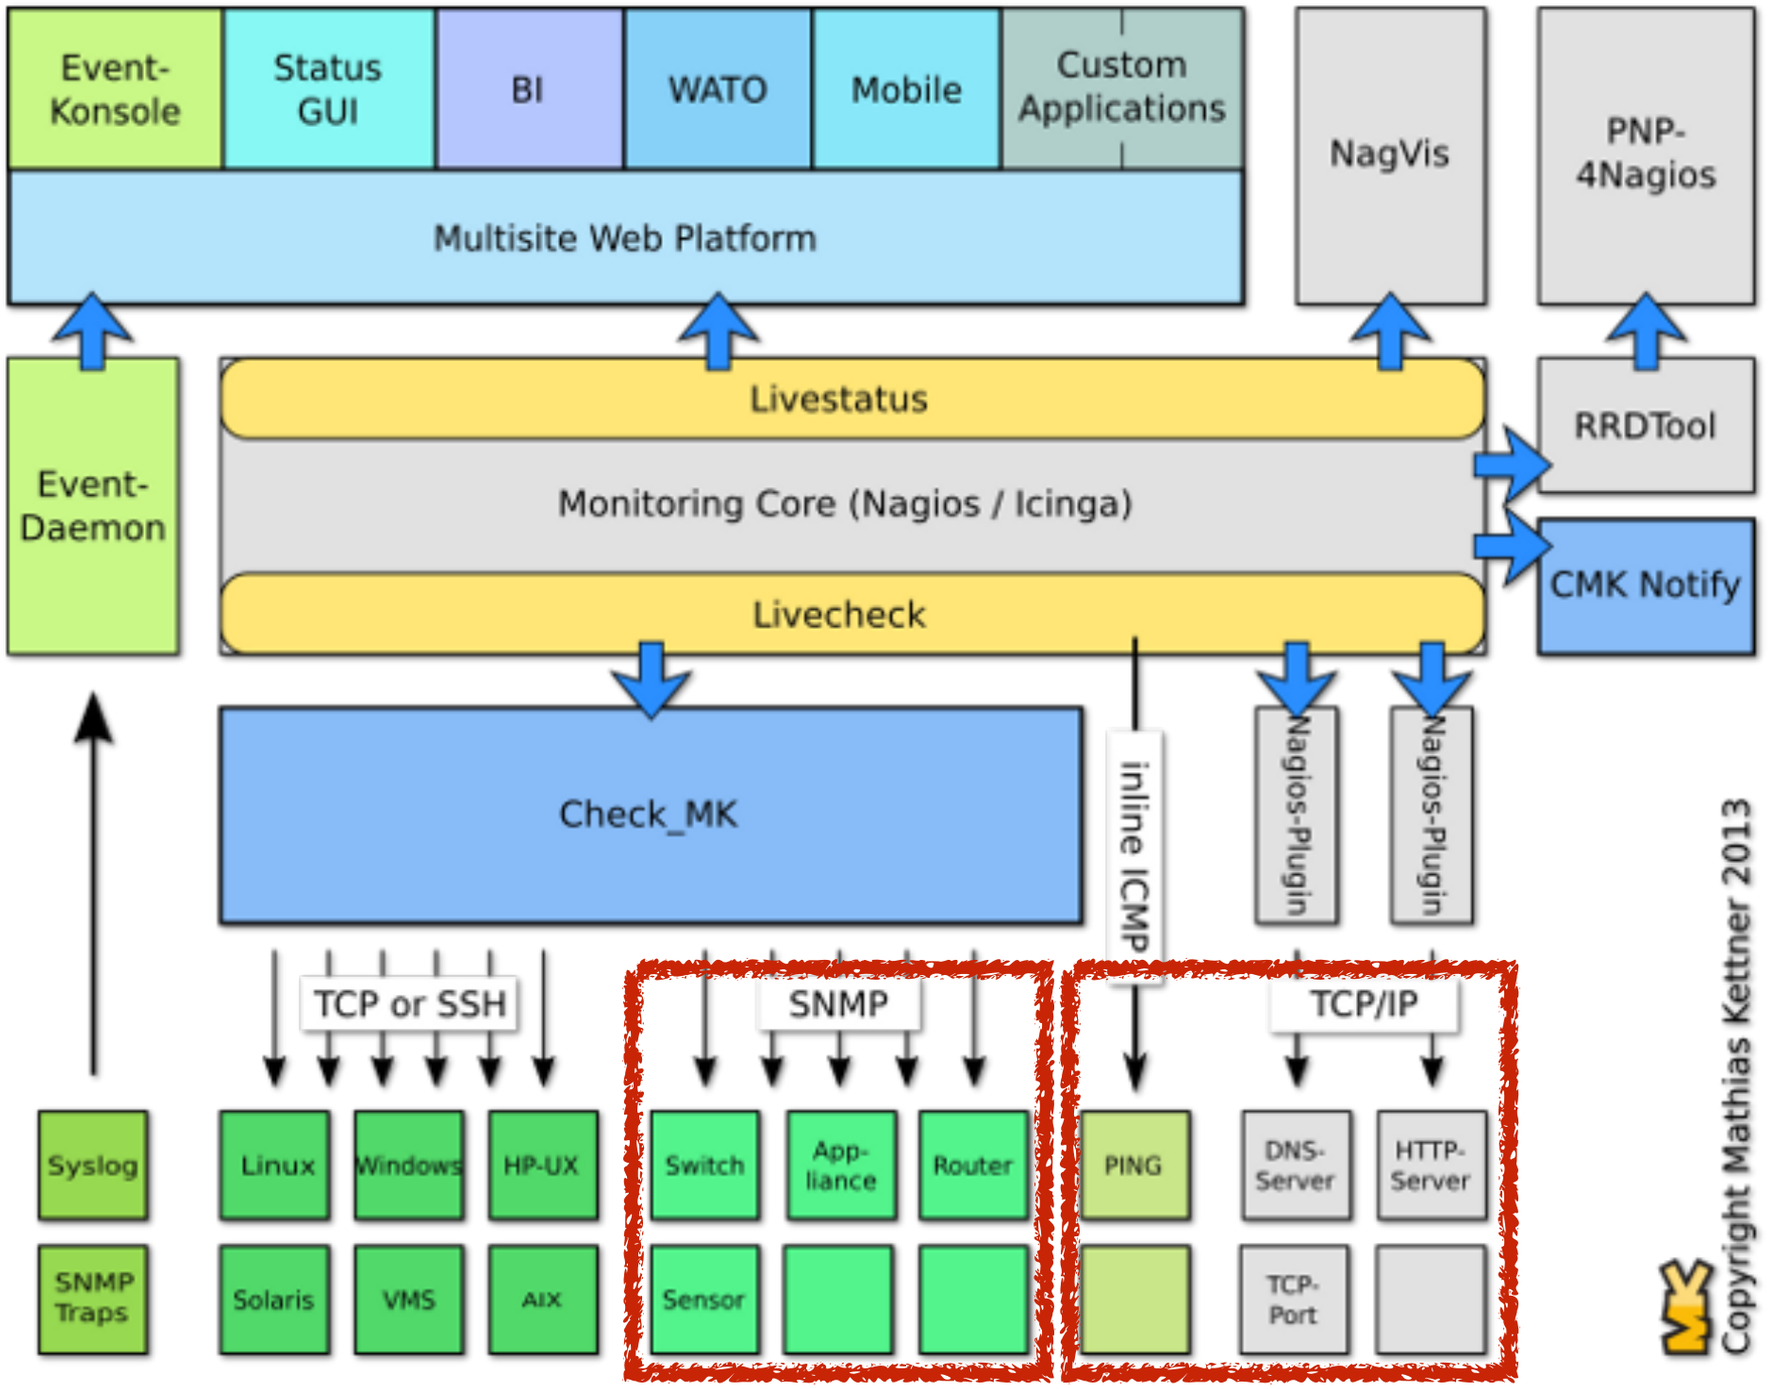

# SNMP

- Basic measurements
  - Bytes in/out, interface status, interface speed
- Ability to set thresholds
  - “if the interface goes above X Bps for Y secs then…”
- Visibility on the troubles but…
  - What is the cause?
  - Who is the bad guy?

# Active Measurements

- PINGs, HTTP(s) requests, DNS queries
- Ptolemaic view
  - Measurements depend on the observation point
  - Are the measurements really representative of users' quality of experience (QoE)?

# What is missing today?

- Network traffic visibility
- Device-to-traffic binding
- Distributed service availability and performance

- Network traffic visibility to pretend to be the end-user
- Measurements that are representative of the users' QoE

# Available Monitoring Technologies

- Fortunately, several technologies come into play when it comes to provide visibility into the network traffic

- Switches: sFlow

- Routers: NetFlow v5/v9, IPFIX, NetFlow Lite

# Monitoring From Scratch

- Sometimes it is necessary to start the monitoring from raw packets, for example when
  - NetFlow/sFlow is not available
  - Custom measurements are needed (eg., RTT, Network Latency, DPI)
- An extra piece of software (a **probe**) is required to process the packets and translate them into something actionable
- The **probe** can be fed with packets from
  - Switch mirror ports
  - Network TAPs

# Compressing Raw Packets

- All the monitoring technologies available are inherently connected by the necessity to “compress” packets into actionable summaries that preserve the basic properties of the network communications
  - Often impractical to work with raw network packets
- Network packets are still important for providing evidence or troubleshooting problems (“pcap or it didn't happen!”) but they are “too raw” and take too much storage space.
- Network flow analysis is a good way to “compress packets”: sFlow do it with sampling, NetFlow with stateful connection-based packet classification.

# Network Flows: What Are They?

- “A flow is a set of packets with a set of common packet properties” (e.g. common IP address/port).

- All the packets of a web session can be summarized in a flow
  - “**host 1.2.3.4** fetched website **www.ntop.org** served by **host 6.7.8.9** in **S** seconds [with network latency X ms [and application latency Y ms [and …]]]”

- Other examples of network flows are a Skype/VoIP call, an FTP file transfer, an SSH session, etc.

# Packets vs Flow

| No. | Time | Source | Destination | Source P | Destination P | Protocol | Length | Host | Info |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 10:55:45.533410 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 78 | | 62241 → 80 [SYN] Seq=0 Wi… |
| 2 | 10:55:45.592083 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | TCP | 78 | | 80 → 62241 [SYN, ACK] Seq… |
| 3 | 10:55:45.592149 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [ACK] Seq=1 Ac… |
| 4 | 10:55:45.592321 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | HTTP | 626 | www.magaz… | GET /schluesselanhaenger-… |
| 5 | 10:55:45.753263 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | TCP | 66 | | 80 → 62241 [ACK] Seq=1 Ac… |
| 6 | 10:55:45.859401 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | HTTP/1.1 200 OK [Unreasse… |
| 7 | 10:55:45.860471 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | Continuation |
| 8 | 10:55:45.860540 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [ACK] Seq=561 … |
| 9 | 10:55:45.861612 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | Continuation |
| 10 | 10:55:45.861679 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [ACK] Seq=561 … |
| 11 | 10:55:45.918319 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | Continuation |
| 12 | 10:55:45.918427 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [ACK] Seq=561 … |
| 13 | 10:55:45.919538 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | Continuation |
| 14 | 10:55:45.920543 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | Continuation |
| 15 | 10:55:45.920606 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [ACK] Seq=561 … |
| 16 | 10:55:45.921750 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | Continuation |
| 17 | 10:55:45.921963 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [ACK] Seq=561 … |
| 18 | 10:55:45.922715 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | Continuation |
| 19 | 10:55:45.924202 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | Continuation |
| 20 | 10:55:45.924276 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [ACK] Seq=561 … |
| 21 | 10:55:45.976788 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 1506 | | Continuation |
| 22 | 10:55:45.977014 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | HTTP | 598 | | Continuation |
| 23 | 10:55:45.977153 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [ACK] Seq=561 … |
| 24 | 10:56:00.979471 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | TCP | 66 | | 80 → 62241 [FIN, ACK] Seq… |
| 25 | 10:56:00.979522 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [ACK] Seq=561 … |
| 26 | 10:56:01.007866 | 192.168.1.110 | 212.1.42.233 | 62241 | 80 | TCP | 66 | | 62241 → 80 [FIN, ACK] Seq… |
| 27 | 10:56:01.064432 | 212.1.42.233 | 192.168.1.110 | 80 | 62241 | TCP | 66 | | 80 → 62241 [ACK] Seq=1493… |

| | Application | L4 Proto | Client | Server | Duration | Breakdown | Actual Thpt | Total Bytes | Info |
|---|---|---|---|---|---|---|---|---|---|
| Info | HTTP | TCP | 192.168.1.110:62241 | www.magazin.com :http | 16 sec | Server | 0 bit/s | 16.89 KB | /schluesselanhaenger-key... |

# SNMP vs Flow

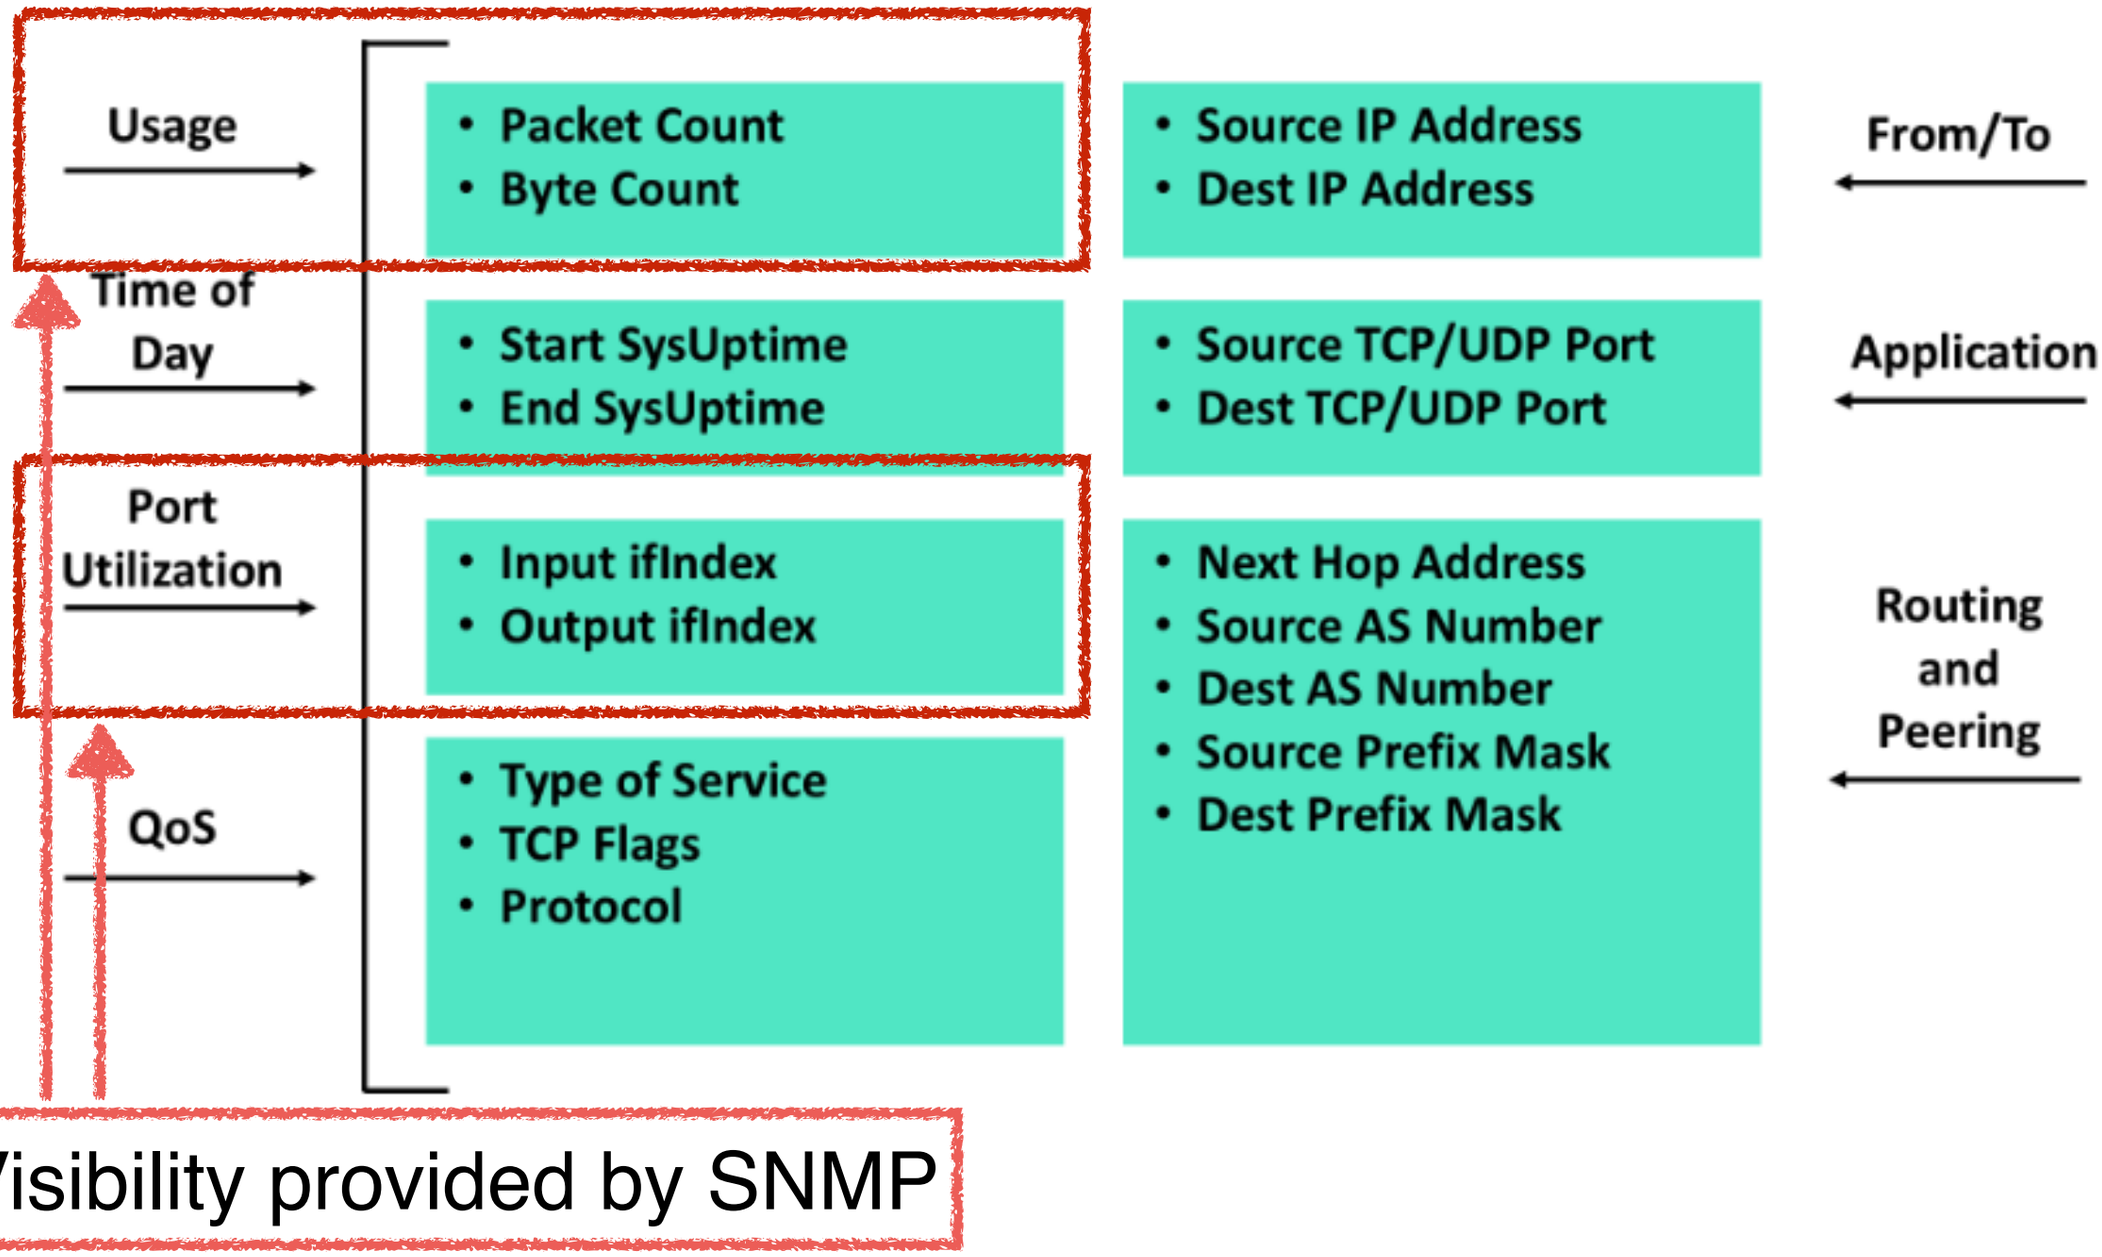

# What to Expect from Flows

- PC/services that are heavily using the network (Top Talkers) — Who's making the network slow?

- Highlight the sources and destinations of network traffic — There's a file transfer going on to a Chinese host

- Application protocols used (Skype, HTTP, Email) — Someone is watching Netflix at work!

- Advanced reporting (billing and accounting) — What was the amount of bandwidth consumed over the past week?

- Legitimate but unauthorized/suspicious traffic (eg. Tor or VPN)

# What NOT to Expect from Flows

- Non-IP traffic (e.g. NetBIOS, AppleTalk).
- L2 information (e.g. interface up/down state changes)
- Filtered traffic (e.g. firewall policy counters).
- Per-link statistics (e.g. link usage, congestion, delay, packet loss).

# What Pieces are Needed?

- A collector
  - To speak sFlow/NetFlow with switches and routers
- A probe
  - To generate flows out of network TAPs and port mirrors
- Least but not last, a good tool for visualization and analysis…

# nProbe

- NetFlow v5/v9/IPFIX, NetFlow-Lite and sFlow collector
- 10Gbps+ probe with DPI
- Extensible (support plugins)
- Convert flow format (sFlow-to-NetFlow/IPFIX) or version (e.g. v5 to v9)
- Ability to export to Kafka, MySQL, ElasticSearch, Text Files, Syslog, JSON, ZMQ, …

# Collector vs Probe Mode

Probe

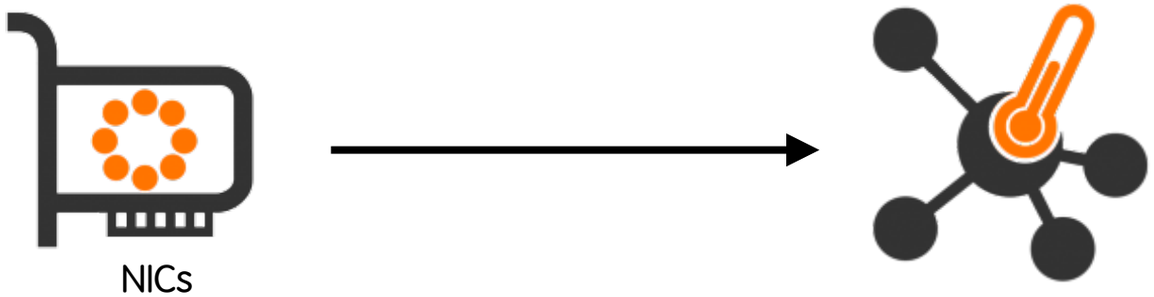


Collector

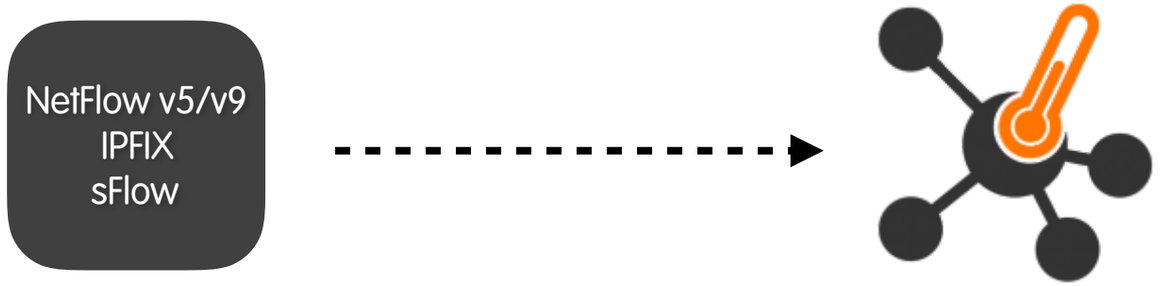

# Deep Packet Inspection with nDPI

- nProbe in collector mode performs Deep Packet Inspection (DPI) using the opensource library nDPI
- Supported protocols (> 240) include:
  - P2P (Skype, BitTorrent)
  - Messaging (Viber, Whatsapp, MSN, The Facebook)
  - Multimedia (YouTube, Last.gm, iTunes)
  - Conferencing (Webex, CitrixOnLine)
  - Streaming (Zattoo, Icecast, Shoutcast, Netflix)
  - Business (VNC, RDP, Citrix, *SQL)

# nProbe Downstream Export

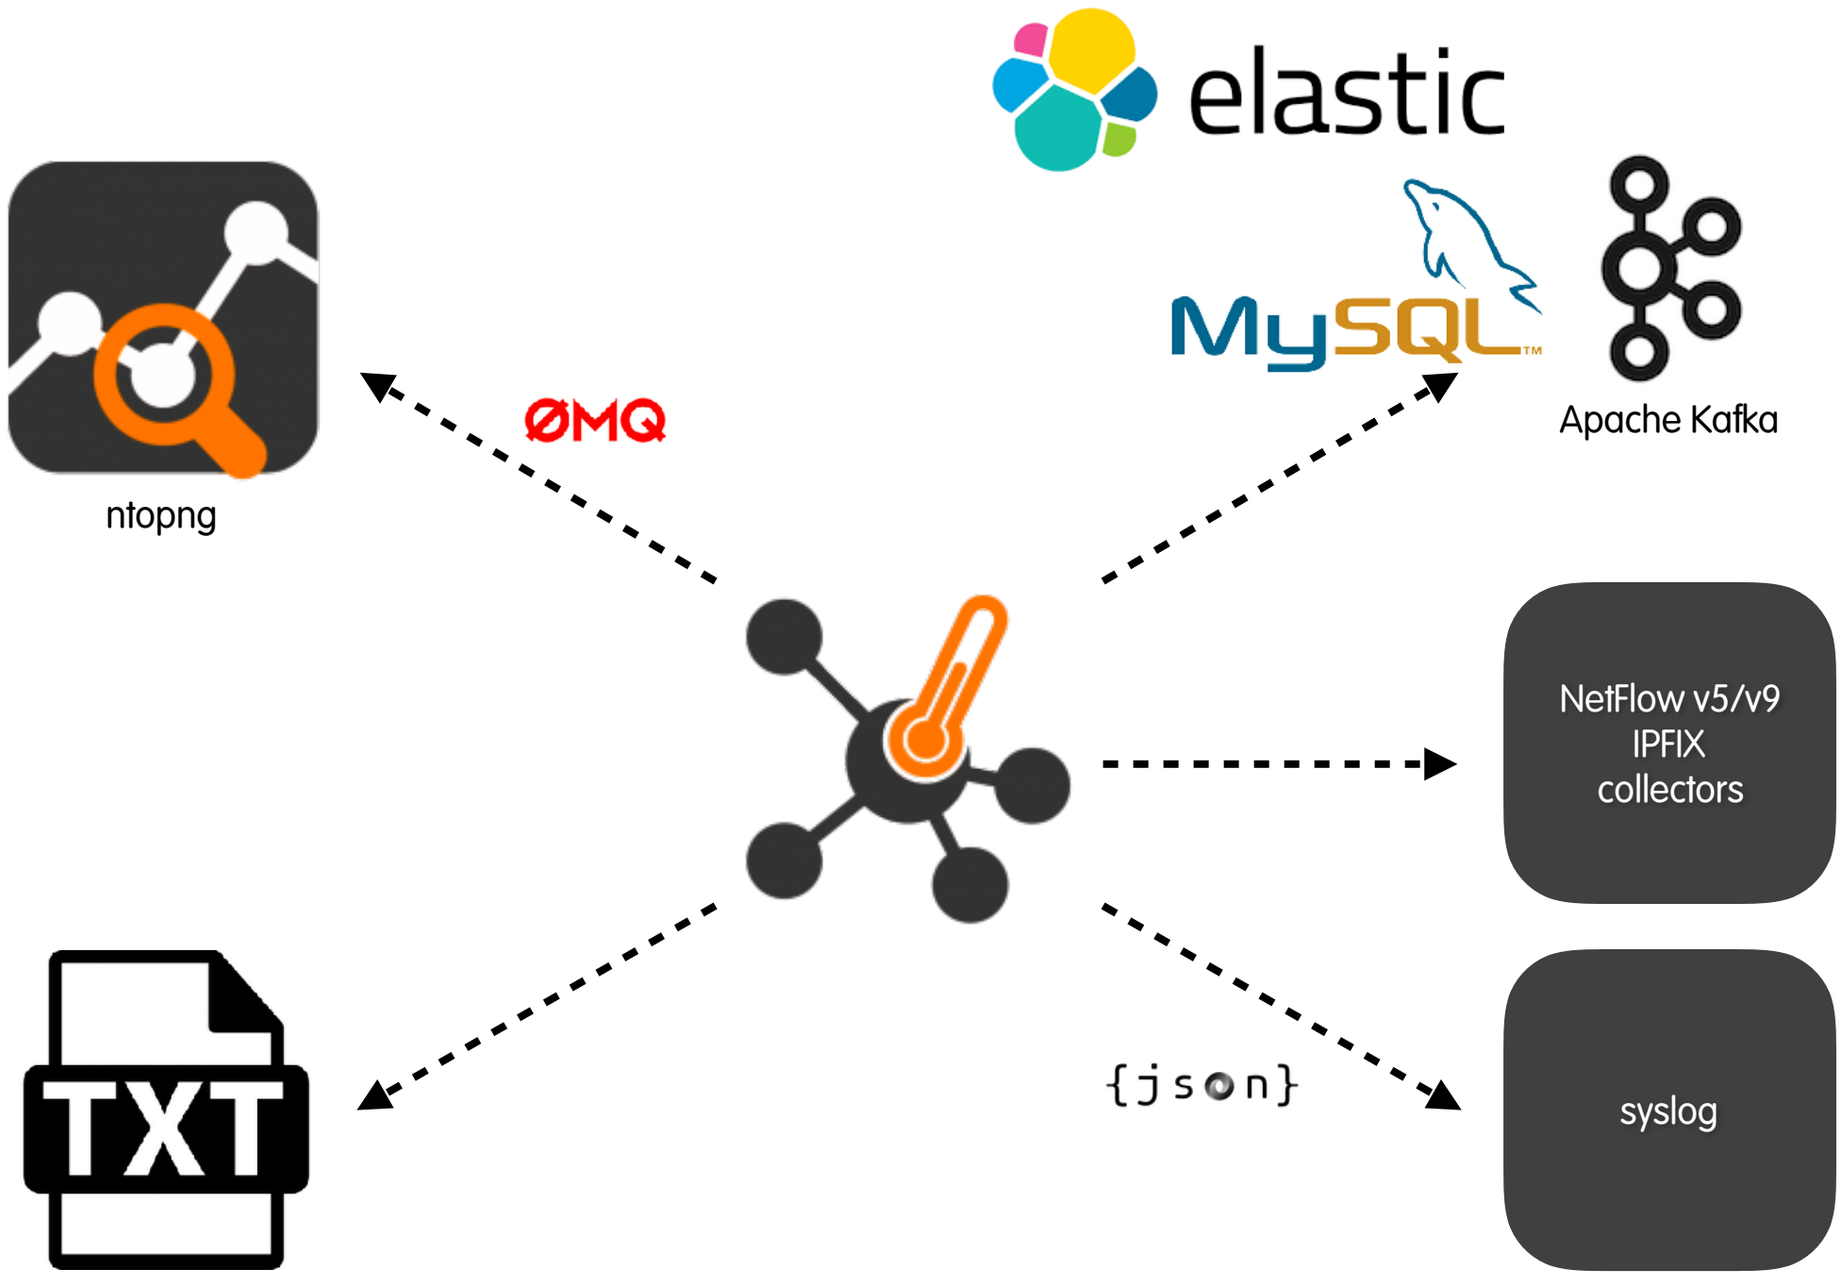

# Flow Analysis with ntopng

- ntopng is a monitoring tool capable of harvesting nProbe-generated data for visualization and analysis

- Web-based HTML5 Graphical User Interface

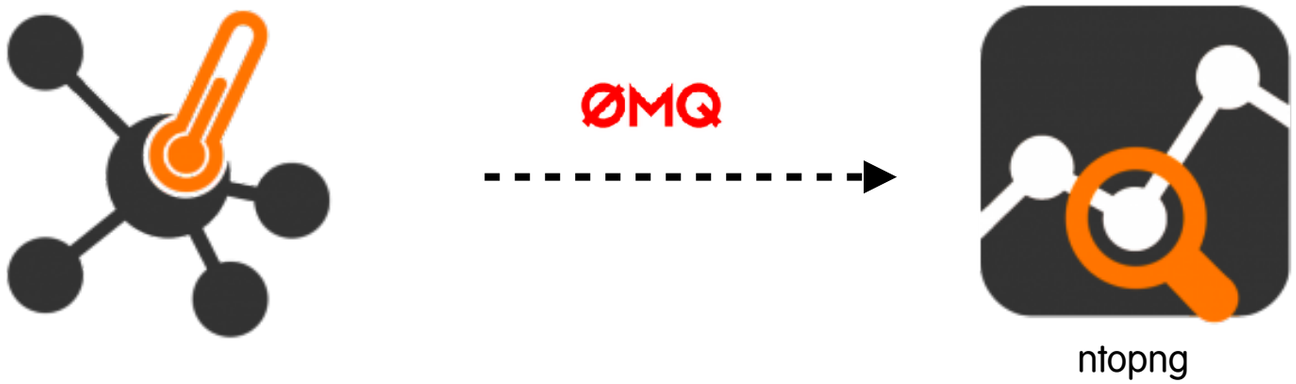

# Main ntopng Features

- Embedded alerting system with several external endpoints including nagios, email and slack
- Grafana datasource
- Ready for
- Support for NetFlow/sFlow/SNMP
- Passive/Active Network Device Discovery
- Traffic Behavior Analysis

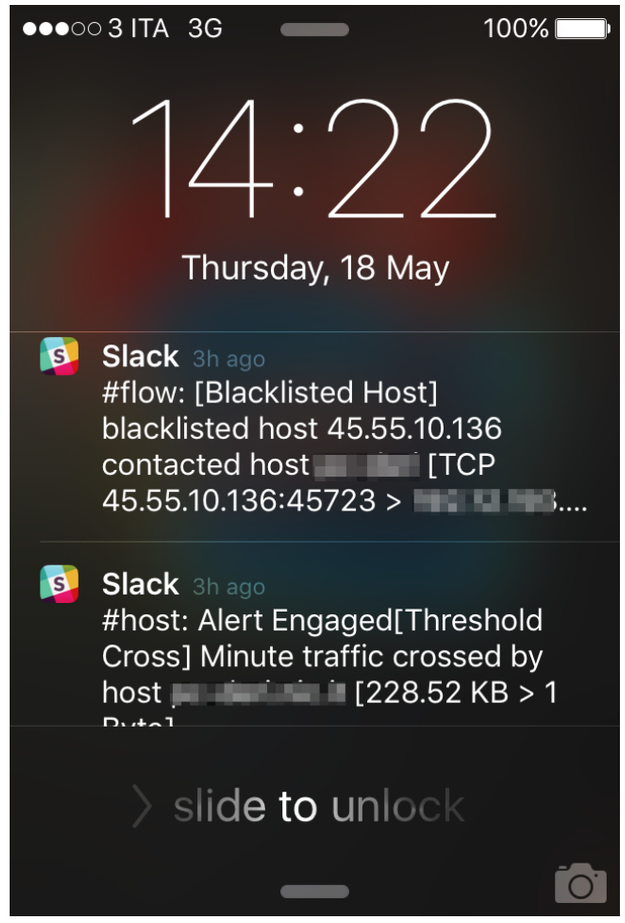

# ntopng Dashboard

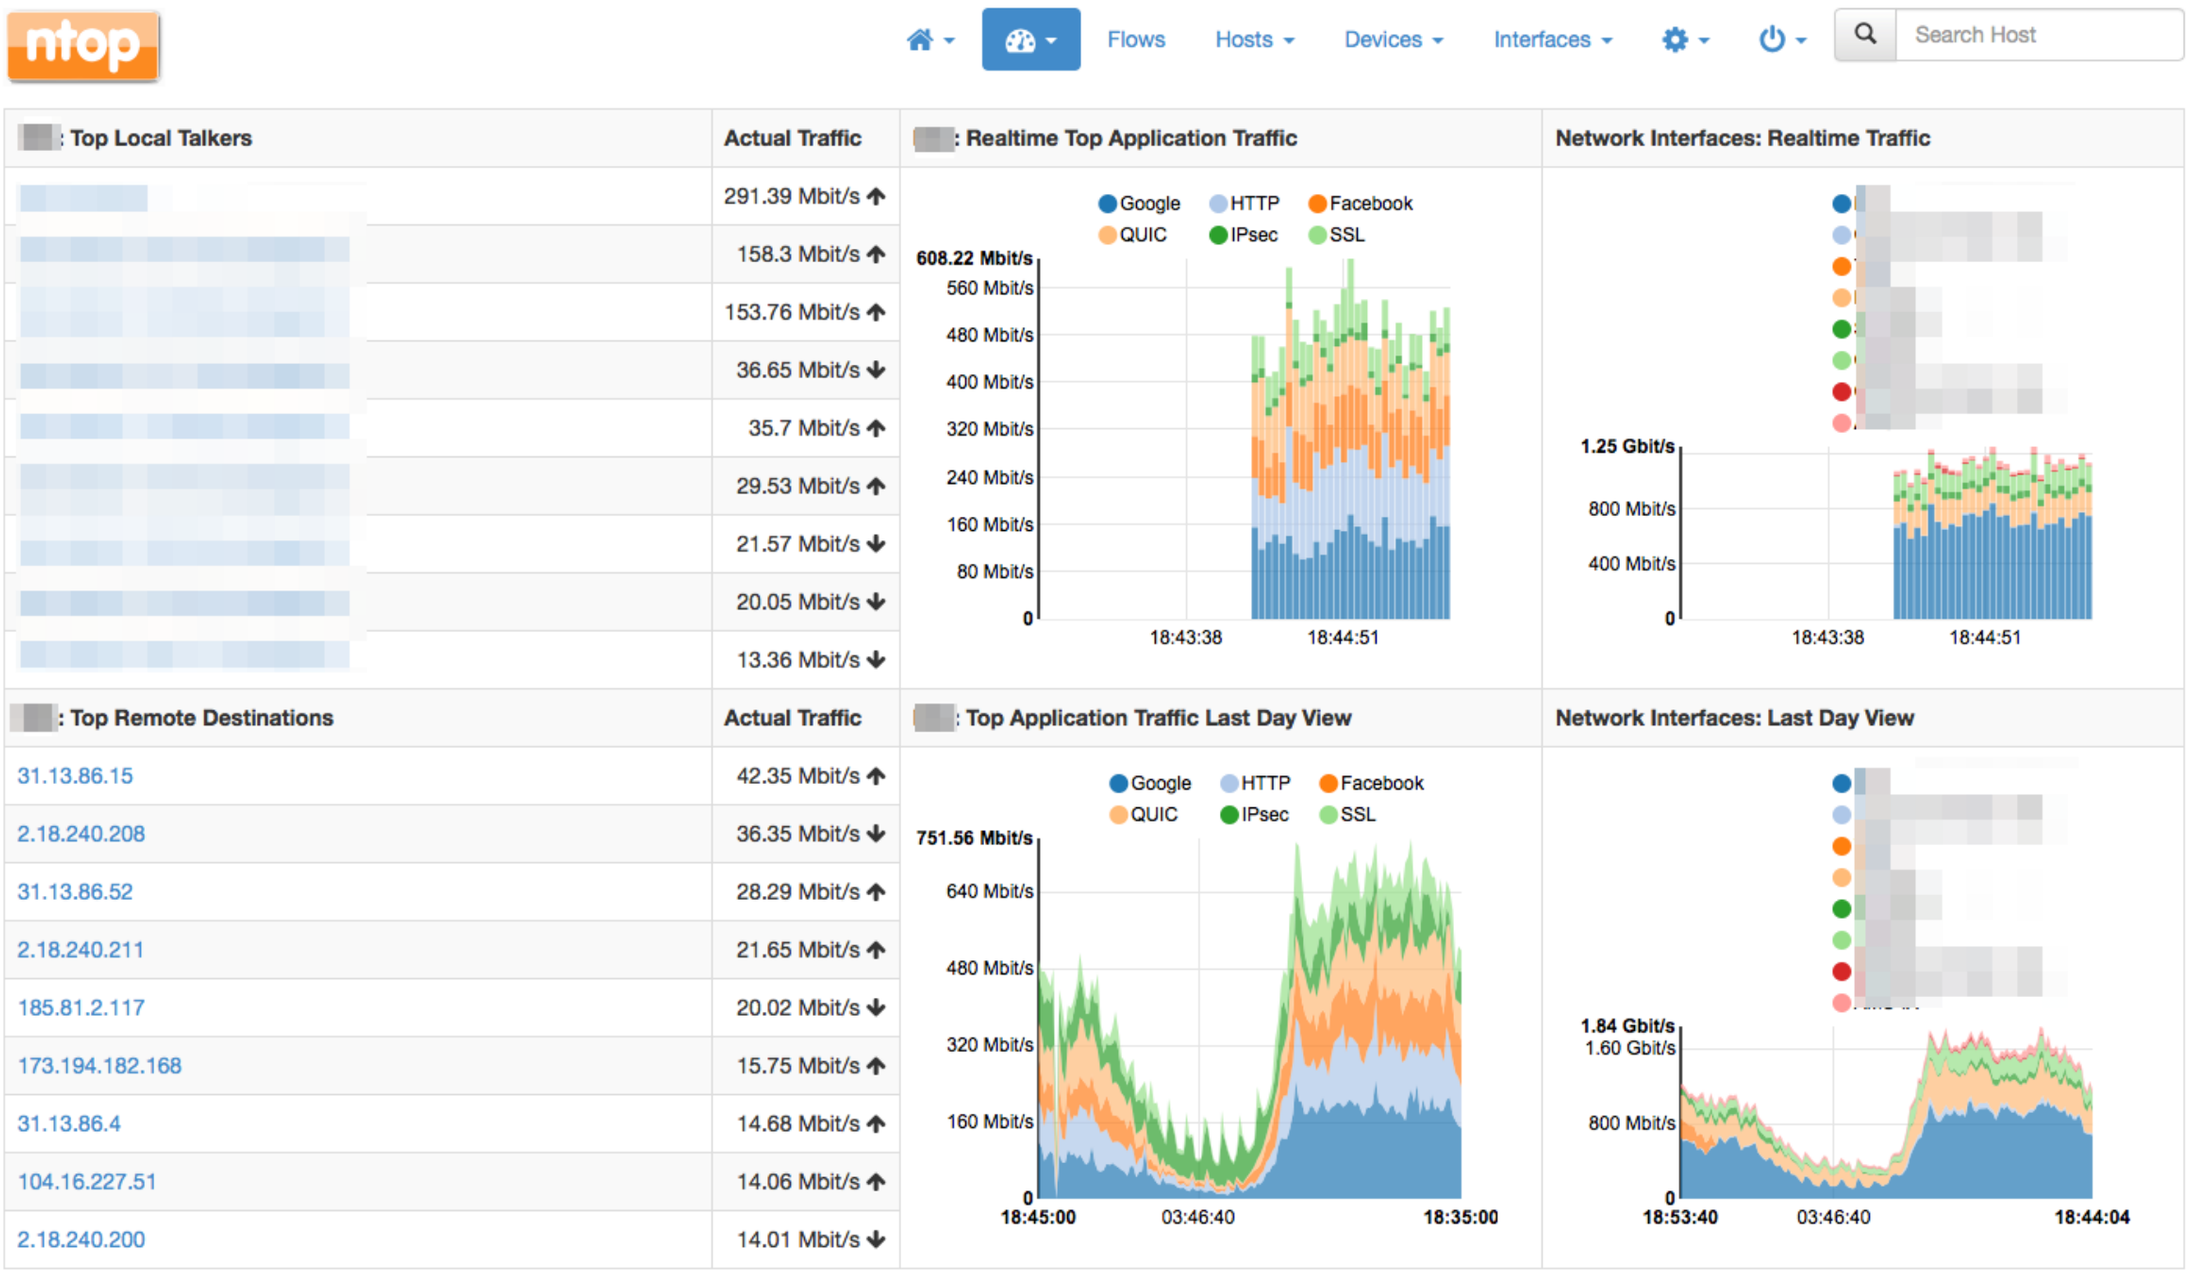

# SNMP Monitoring

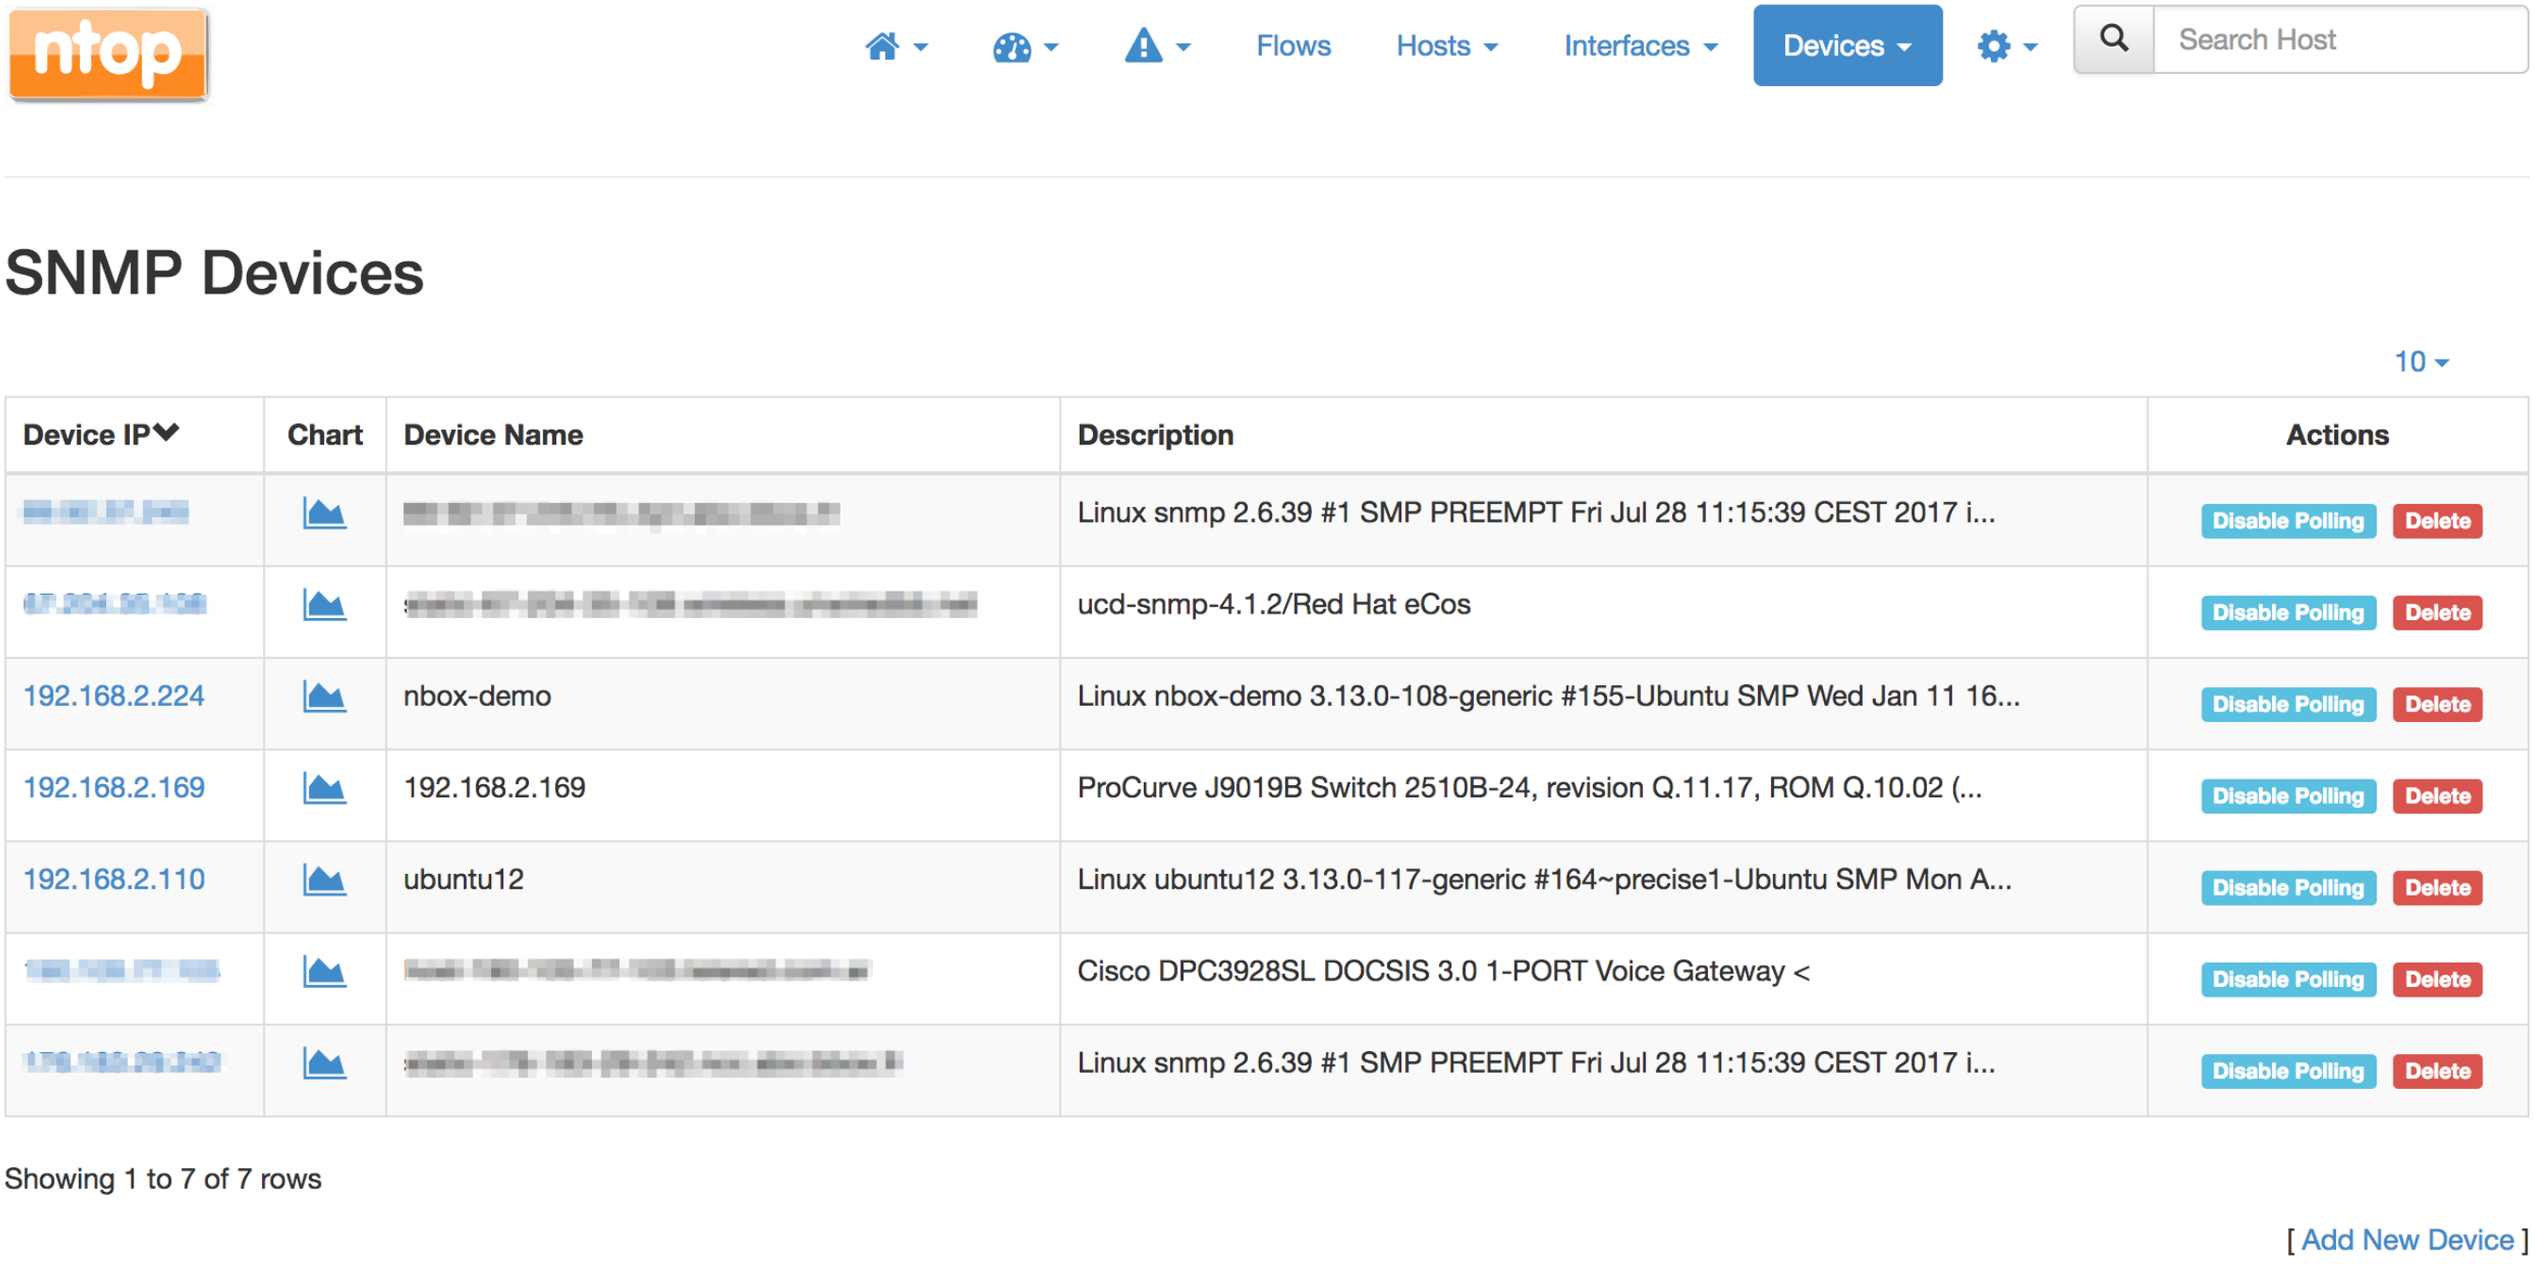

# SNMP Interfaces Slicing & Dicing


ntop

Flows Hosts Interfaces Devices Search Host

SNMP Device: 192.168.2.169 Interfaces

## Interfaces

10 Interface Type Status

| | Index | Name | VLAN | Speed | Status | In Bytes | Out Bytes | In Discar | ge |
|---|---|---|---|---|---|---|---|---|---|
| Info | 25 | 25 | 1 | 1 Gbit | Up | 88.18 GB | 1001.33 GB | | nin, 11 sec |
| Info | 9 | 9 | 1 | 100 Mbit | Up | 1.04 TB | 42.63 GB | | nin, 20 sec |
| Info | 15 | 15 | 1 | 100 Mbit | Up | 2.28 GB | 29.94 GB | | 13 days, 13 min, 53 sec |
| Info | 19 | 19 | 1 | 100 Mbit | Up | 1.71 GB | 26.23 GB | | 3 h, 24 min, 2 sec |
| Info | 21 | 21 | 1 | 100 Mbit | Up | 340.65 MB | 9.26 GB | | 1 h, 4 min, 52 sec |
| Info | 8 | 8 | 1 | 100 Mbit | Up | 167.48 MB | 635.18 MB | | 55 days, 7 h, 17 min, 13 sec |
| Info | 4152 | HP ProCurve Switch software loopback interface | | | Up | 0 Bytes | 0 Bytes | | |
| Info | 57 | DEFAULT_VLAN | | | Up | 0 Bytes | 0 Bytes | | 55 days, 7 h, 17 min, 11 sec |

All Types
softwareLoopback
ethernetCsmacd
propVirtual

Showing 1 to 8 of 8 rows

# SNMP and Flow Collection: Connecting the Dots [1/2]

- ntopng for each flow exporter device is able to detect if there is a corresponding SNMP device configured and glue them up.

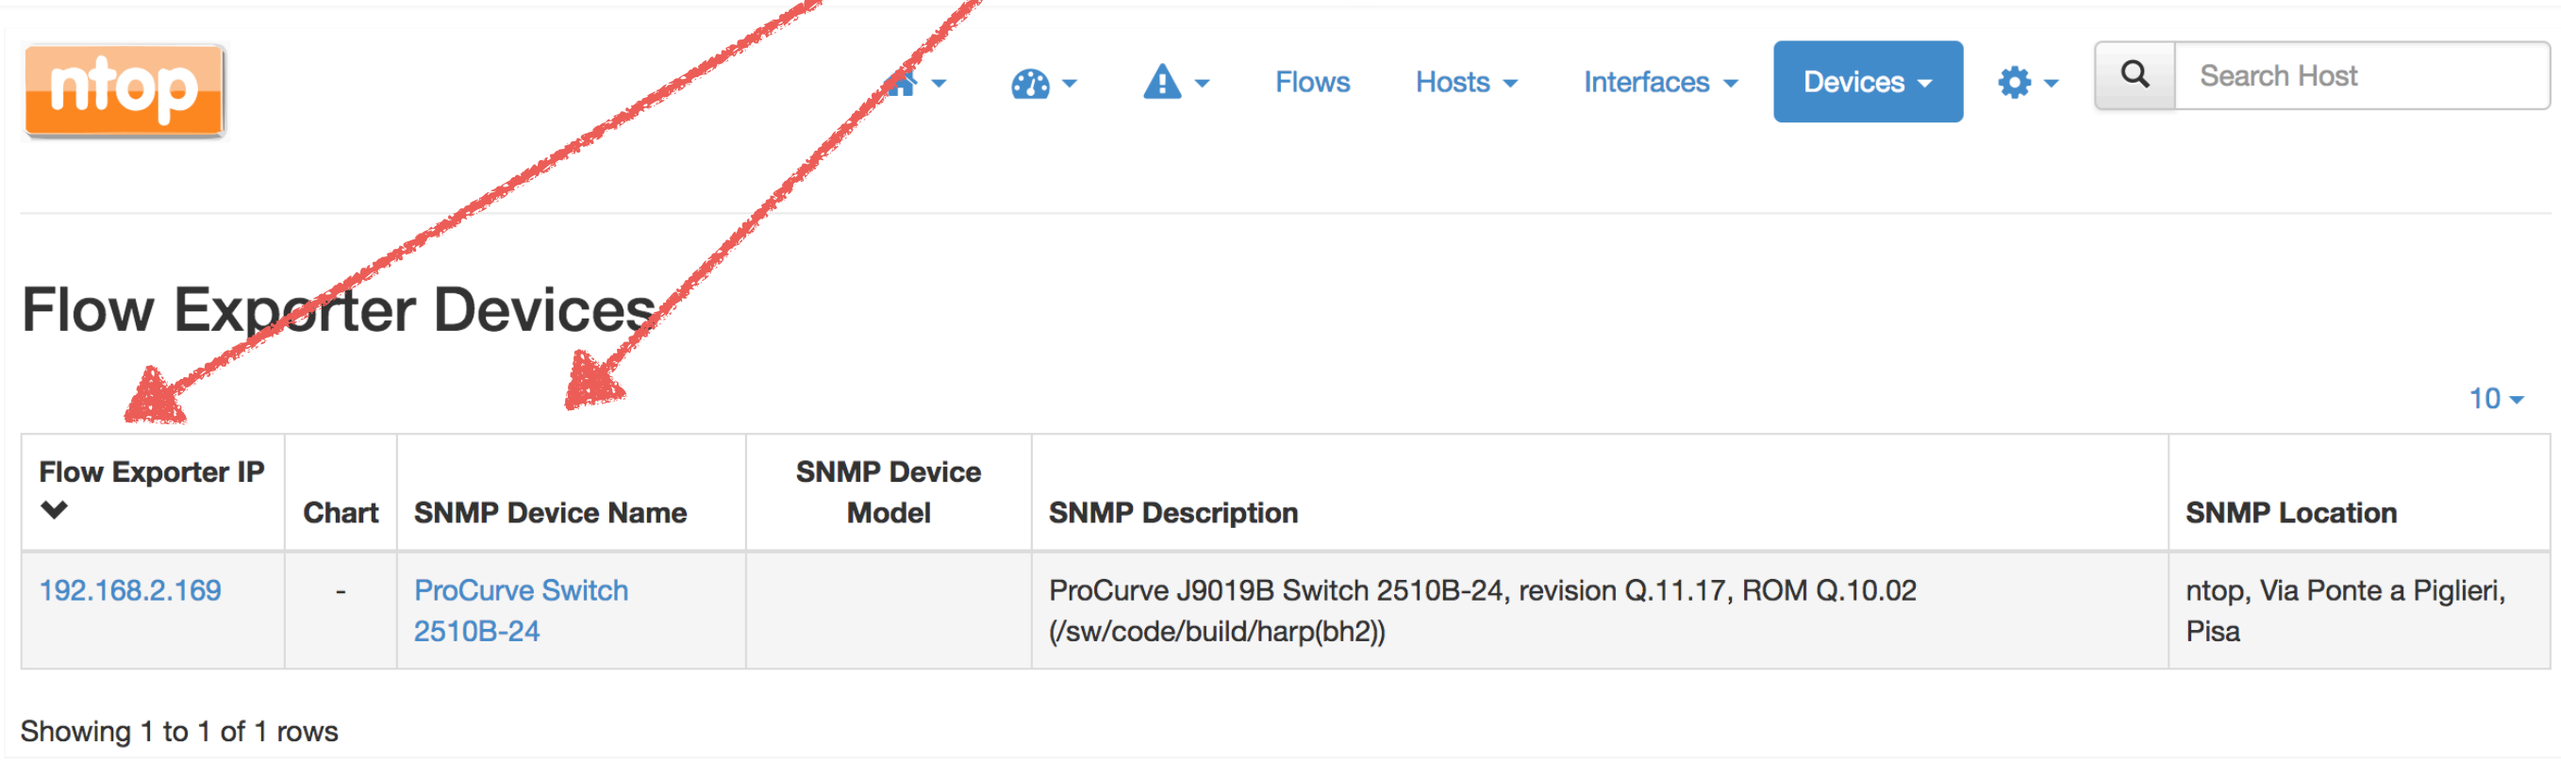

# SNMP and Flow Collection: Connecting the Dots [2/2]


ntop

Flows | Hosts | Interfaces | Devices | Search Host

## Recently Active Flows [Flow Exporter 192.168.2.169]

10 | Hosts | Status | Direction | Applications | Categories | Input Interface | Output Interface | Flow Exporter | IP Version

All Flow Exporters

Flow Exporter 192.168.2.169[ProCurve Switch 2510B-24...]

| | Application | L4 Proto | Client | Server | Duratic | | | | nfo |
|---|---|---|---|---|---|---|---|---|---|
| Info | SSH | TCP | ubuntu:54857 | devel:ssh | 12 sec | | | | |
| Info | UPnP | UDP | 192.168.2.136:3116 | 239.255.255.250:1900 | 0 sec | Client | 253.14 bit/s ↑ | 1.2 KB | |
| Info | DNS | UDP | devel:35640 | 199.85.126.30 :domain | 0 sec | Clier Server | 51.07 bit/s ↑ | 236 Bytes | |
| Info | DNS | UDP | devel:26111 | 199.85.126.30 :domain | 0 sec | Client Server | 57.81 bit/s ↑ | 231 Bytes | |
| Info | DNS | UDP | devel:24017 | 199.85.126.30 :domain | 0 sec | Clien Server | 48.48 bit/s ↑ | 224 Bytes | |
| Info | DNS | UDP | devel:44978 | 199.85.127.30 :domain | 0 sec | Client Server | 40.86 bit/s ↑ | 199 Bytes | |
| Info | DNS | UDP | devel:44978 | 199.85.126.30 :domain | 0 sec | Client Server | 40.86 bit/s ↑ | 199 Bytes | |
| Info | DNS | UDP | devel:13005 | 199.85.126.30 :domain | 0 sec | Client Server | 42.71 bit/s ↑ | 192 Bytes | |
| Info | ICMP | ICMP | devel | 199.85.127.30 | 0 sec | Client | 31.82 bit/s ↑ | 155 Bytes | Echo Reply |
| Info | DNS | UDP | devel:50869 | 199.85.126.30 :domain | 0 sec | Client Server | 29.68 bit/s ↑ | 152 Bytes | |

# Know What's Wrong: Alerts [1/3]

Extensible Thresholds on Traffic/SNMP/NetFlow…

Time Period

Host: 192.168.2.20 | Traffic | Packets | Ports | Peers | Protocols | DNS | HTTP | Flows | SNMP | Talkers

Every Minute | Every 5 Minutes | Hourly | Daily

| Threshold Type | Host Lucas-iMac Thresholds | Local Hosts Common Thresholds |
|---|---|---|
| **Activity Time** Activity time delta (seconds). | > | > |
| **Traffic** Layer 2 bytes delta (sent + received) | > | > |
| **DNS Traffic** Layer 2 bytes delta (sent + received) for DNS detected traffic | > | > |
| **Flows** Flows delta (as client + as server) | > | > |
| **Idle Time** Idle time since last packet seen (seconds) | > | > |
| **P2P Traffic** Layer 2 bytes delta (sent + received) for peer-to-peer detected traffic | > | > |

# Know What's Wrong: Alerts [2/3]

Open Issues

Past Issues

Flow Issues

Engaged Alerts | Past Alerts | Flow Alerts

## Engaged Alerts

Who

10

| Date/Time | Duration | Severity | Alert Type | Description |
|---|---|---|---|---|
| Sat May 6 13:03:03 2017 | 2 min, 4 sec | Error | Threshold Cross | Threshold **active** crossed by host [redacted] [65 > 1] |

Showing 1 to 1 of 1 rows

What

When

How Long

# Know What's Wrong: Alerts [3/3]

- External Alerts Endpoints
- Slack and Email
- Nagios via NSCA client
- Nagios will intercept all alerts that are explicitly declared as passive services

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| ntopng-host | NtopngAlert | ? | OK | 12-23-2015 15:25:50 | 0d 17h 27m 59s | 1/1 | Alert for host Y! |
| | NtopngAlert_192.168.1.15_min_bytes | ? | OK | 12-23-2015 09:13:22 | 0d 6h 47m 34s | 1/1 | OK, alarm deactivated |
| | NtopngAlert_192.168.2.0/24 | ? | OK | 12-23-2015 11:02:34 | 0d 4h 33m 4s | 1/1 | OK, alarm deactivated |
| | NtopngAlert_192.168.70.0/24_min_egress | ? | WARNING | 12-23-2015 15:33:01 | 0d 0h 6m 5s | 1/1 | Threshold egress crossed by network 192.168.70.0/24 [1180 > 10] |
| | NtopngAlert_192.168.70.0/24_min_ingress | ? | WARNING | 12-23-2015 15:33:01 | 0d 0h 2m 5s | 1/1 | Threshold ingress crossed by network 192.168.70.0/24 [11241211 > 10] |

# Take Home

- SNMP is OK but it's better if it can be enriched with network traffic
- Network traffic can be compressed with into meaningful representations called flows
- Flow can be collected from sFlow/NetFlow devices or generated with a network probe
- nProbe
  - 10+ Gbps probe
  - NetFlow v5/v9/IPFIX collector
- ntopng
  - Web-based GUI for visualization and analysis
  - Able to collect monitored traffic from remote nProbes
  - Present and past host activities visualization, including ability to alert on suspicious behaviors


# Thank you!

Simone Mainardi, PhD
@simonemainardi
mainardi@ntop.org