Practical Monitoring with ntopng

Matteo Biscosi

ntop

PacketFest'25

Agenda

- Scenario explaination
- How to start and configure ntopng
- ClickHouse or not?
- Timeseries
- Enchant the Analysis

Scenario

- Needs: monitor an entire network (the size of the network does not matter); also the data needs to be seen even after some days (in case of attacks or network analysis)
- Various switches/routers are present in the network
- All the traffic is forwarded by these devices to a single machine by using NetFlow/sFlow

Scenario

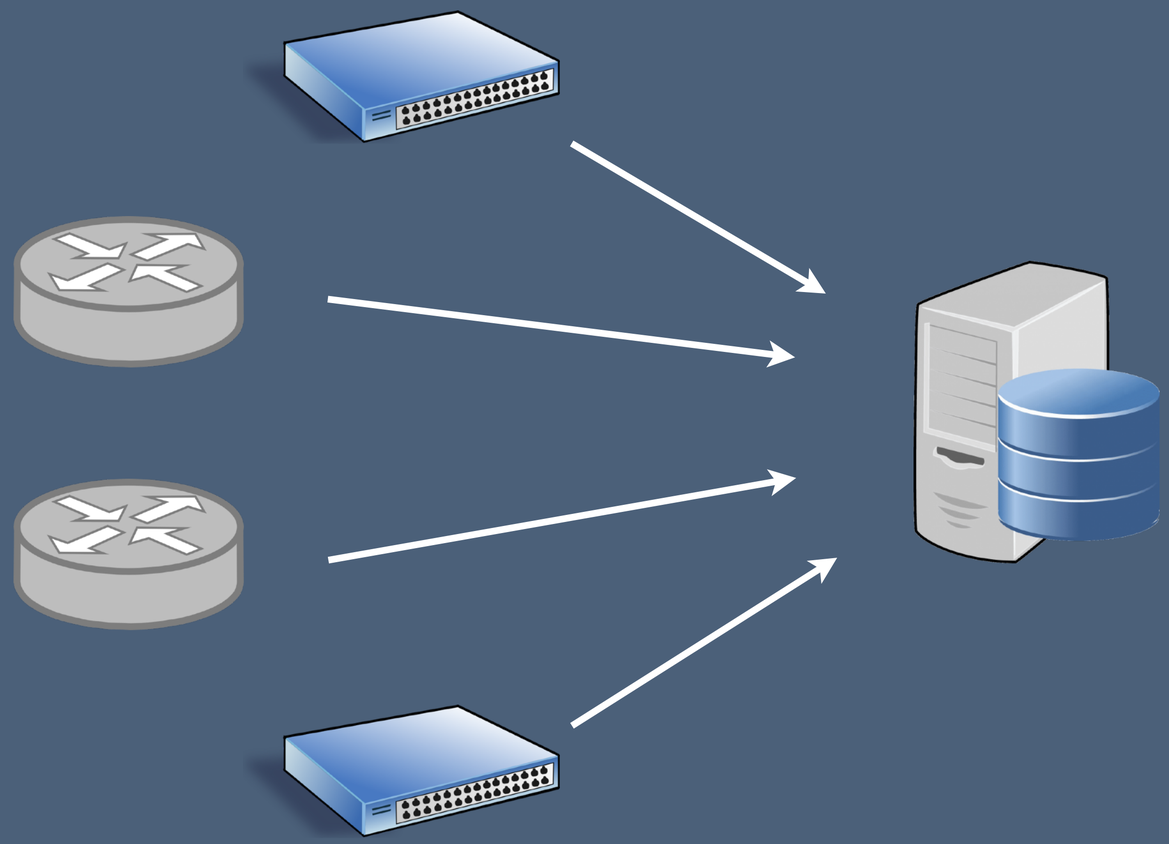

Solution 1

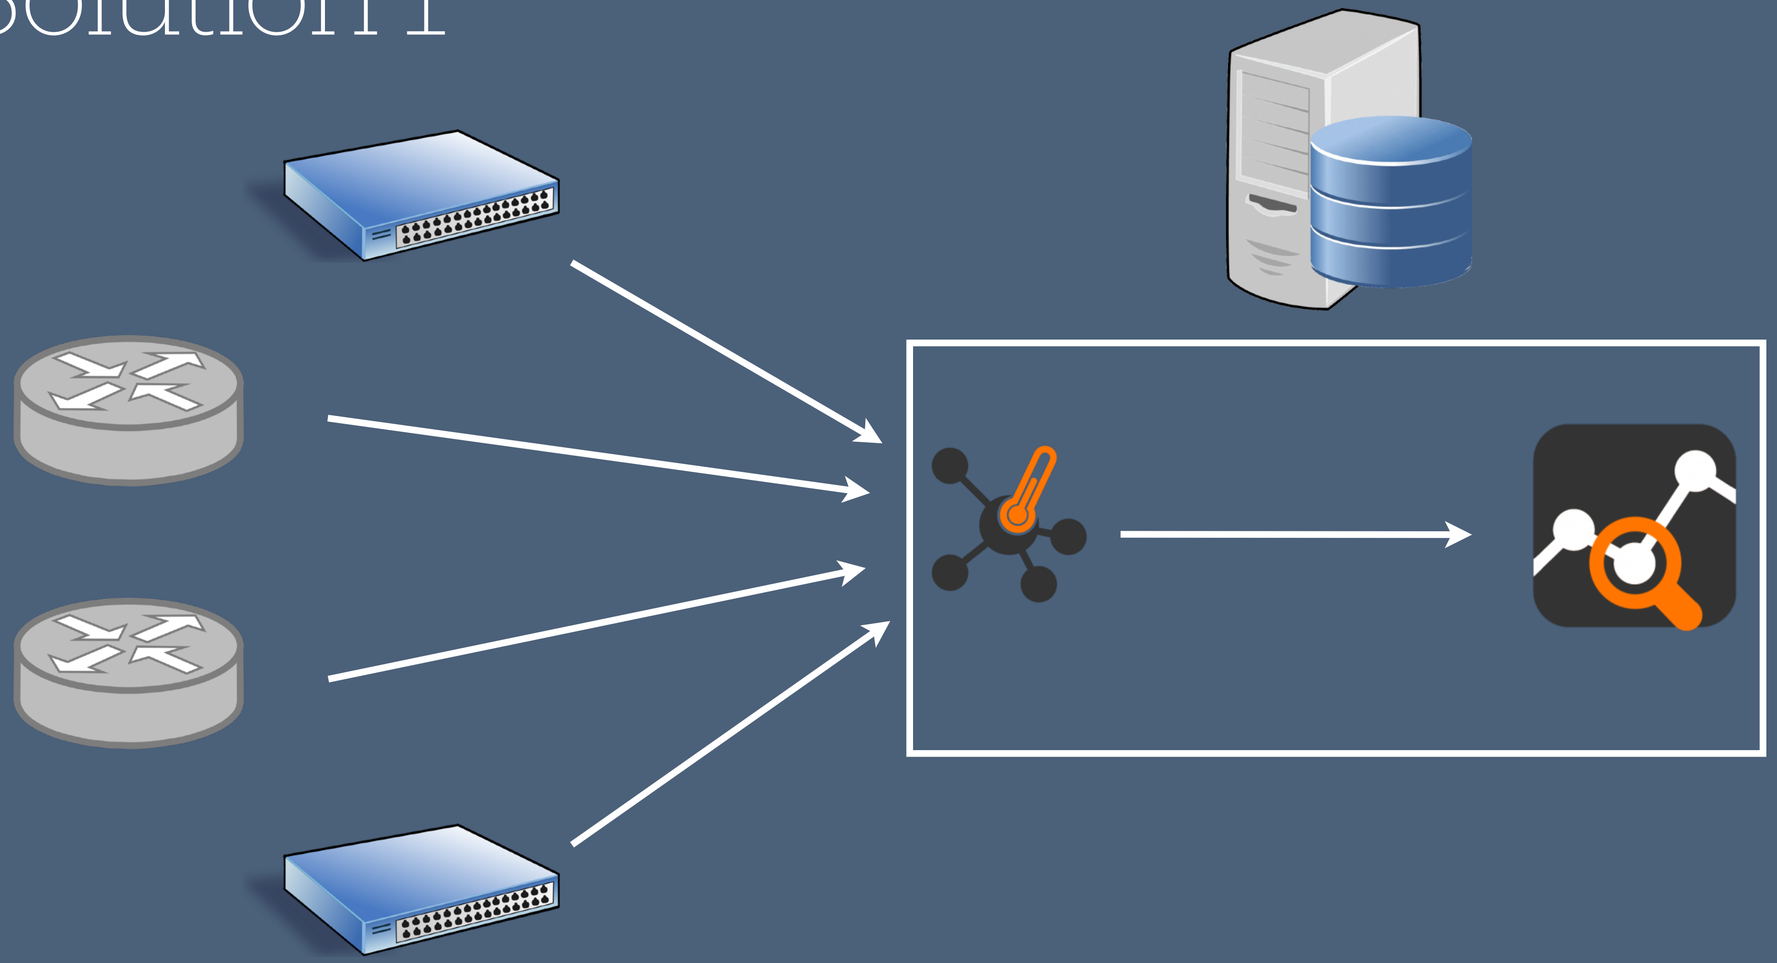

Solution 2

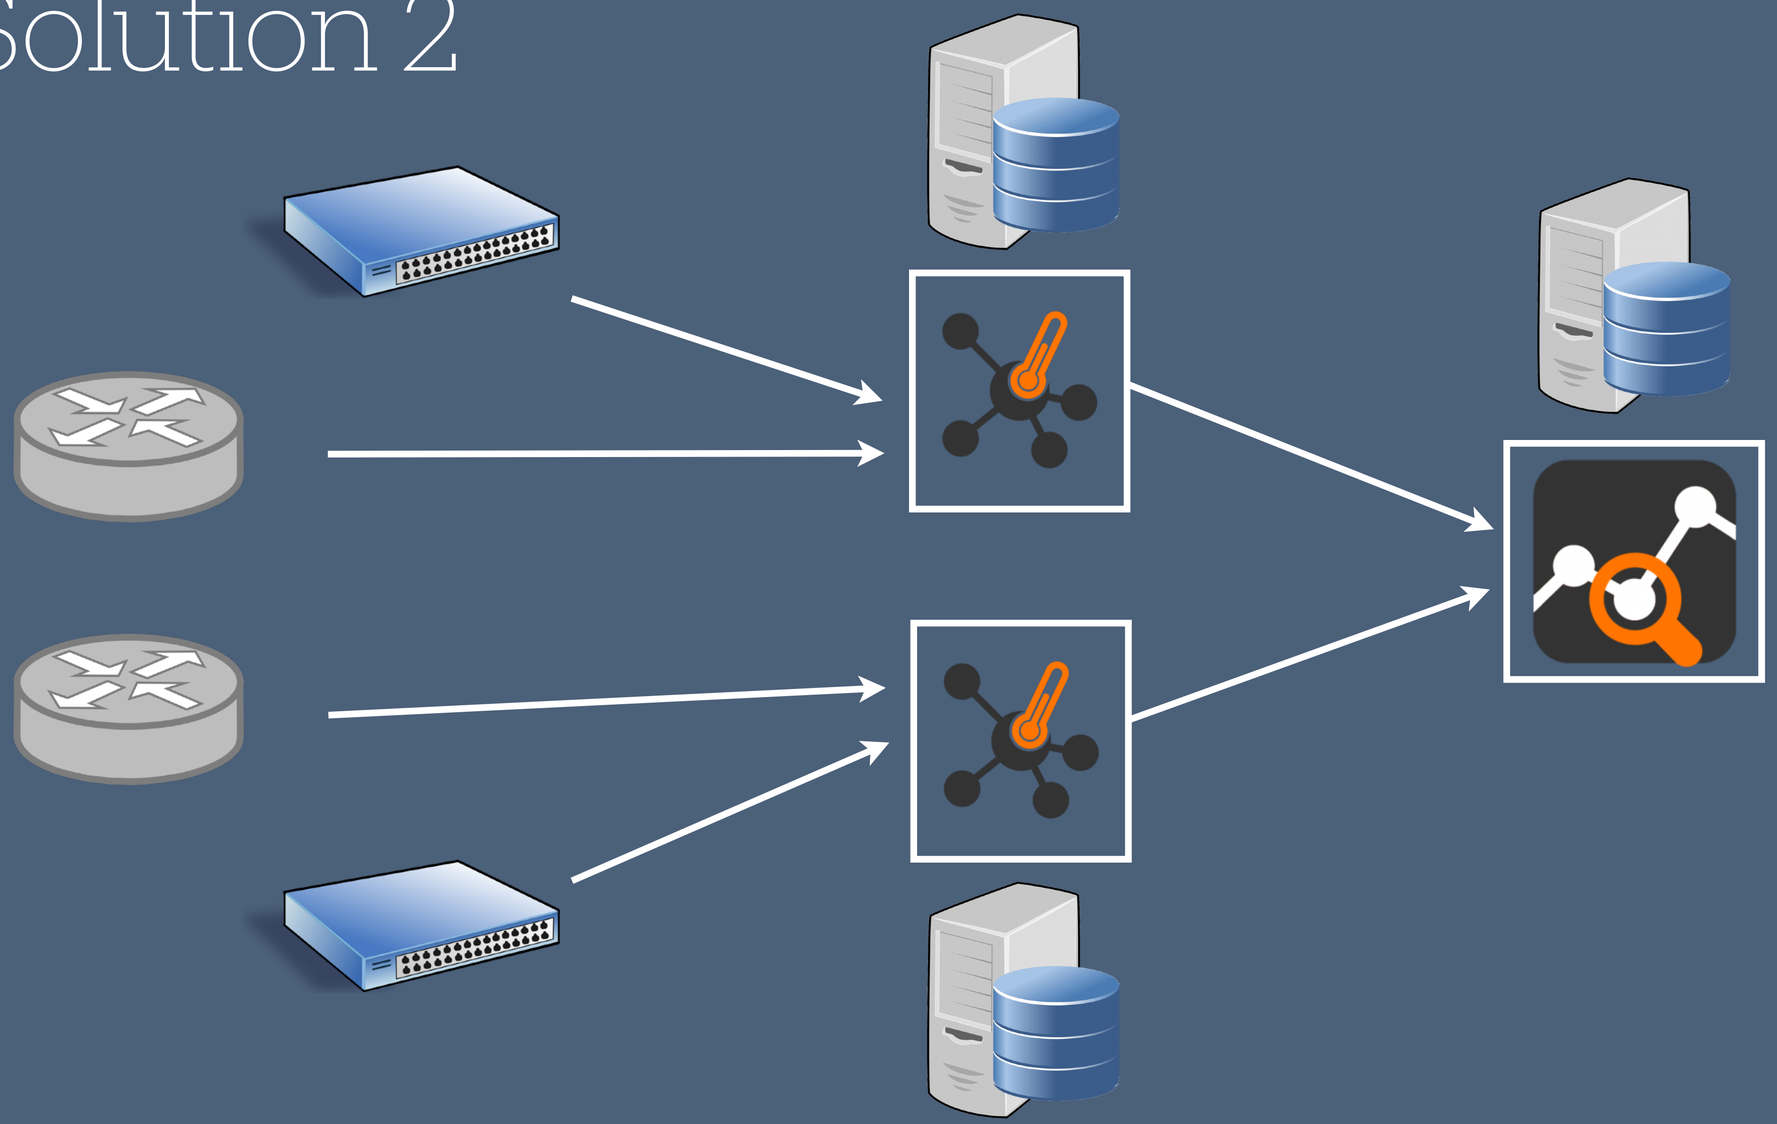

Solutions

- Solution 1:
  - Used in cases the amount of traffic is not too high (depends on the server specifics)
  - In case the network is not heavily segmented
- Solution 2:
  - Used in scenarios where the amount of traffic is quite high (100+ Gbps)
  - The network is quite segmented (e.g. a lot of subnets) and it's requested to have a good 'representation' of the network

Configure the Switches/Routers

- Configure the router to forward in NetFlow (usually more info are provided) or sFlow (less info are provided, used in cases where the amount of traffic is high, in combo with packet sampling) format the traffic

- Forward the traffic towards the server(s) where the probe is

Note: check the port where the switches/routers are sending the NetFlow/ sFlow to

- In case of high traffic send the traffic of each device (switch/router) to a different port

Configure nProbe

- On the receiving server, configure one or more nProbe(s) (one for each port) in order to receive the device traffic
- Configure the nProbe in order to forward the traffic towards ntopng

(Luca will talk more in depth about nProbe configurations and use cases)

Configure nProbe

```
#       This is a sample configuration file for using nProbe in combination with ntopng,
#       exporting flows over ZMQ from nProbe to ntopng.
#       You can enable this configuration by renaming this file to nprobe.conf and restarting
#       the nprobe service.
#
#       -i|--interface
#       Specifies the physical network interface that nProbe will use to perform the
#       monitoring. To disable monitoring from physical interfaces and use nProbe in
#       collector-only mode specify -i=none and use the -3|--collector-port option.
#
-i=none
#
#       -3|--collector-port
#       Specifies the port that is being used by a NetFlow exporter to send NetFlow to nProbe.
#       Multiple NetFlow exporters can simultaneously send data to nProbe using the same port.
#       If you are processing traffic from an interface with -i please comment this option.
#
--collector-port=6363
#
#
#       --zmq
#       Specifies the ZMQ endpoint for delivering flows to ZMQ subscribers.
#
--zmq=tcp://*:5556
```

Configure nProbe

- nProbe configuration file can be found in /etc/nprobe/nprobe.conf
- The configuration file:
  - -i: specify the network interface to capture packets from, in case of flows, the 'none' needs to be specified
  - --collector-port: specify the port where flows are sent (e.g. 6343)
  - --zmq: specify the host:port where the flows are forwarded (--ntopng, in case it is used in combination with ntopng)

Configure ntopng

```
#
#       -i|--interface
#       Specifies  the  network  interface or collector endpoint to be used by ntopng for network
#       monitoring. On Unix you can specify both the interface name  (e.g.  lo)  or  the  numeric
#       interface id as shown by ntopng -h. On Windows you must use the interface number instead.
#       Note that you can specify -i multiple times in order to instruct ntopng to create  multiple interfaces.
-i=tcp://127.0.0.1:5556
#                                                                                                 #
#       -m|--local-networks
#       ntopng determines the ip addresses and netmasks for each active interface. Any traffic on
#       those  networks  is considered local. This parameter allows the user to define additional
#       networks and subnetworks whose traffic is also considered local in  ntopng  reports.  All
#       other hosts are considered remote. If not specified the default is set to 192.168.1.0/24.
#
#       Commas  separate  multiple  network  values.  Both netmask and CIDR notation may be used,
#       even mixed together, for instance "131.114.21.0/24,10.0.0.0/255.0.0.0".
#
-m="10.10.123.0/24=Milan,10.8.124.0/24=Paris,10.7.10.0/24=Rome,10.6.0.0/24=Florence"
#
#       -X|--max-num-flows
#       Set max number of active flows (default: 131072)
#
-X=500000
#
#       -x|--max-num-hosts
#       Set max number of active hosts (default: 131072)
#
-x=500000
```

Configure ntopng

- ntopng configuration file can be found in /etc/ntopng/ntopng.conf
- The configuration file:
  - -i: specify the network interface to capture packets/flows from, for example in case of nProbe, specify the IP:PORT where nprobe is exporting flows to
  - -m: specifies the local networks
  - -x/-X: specifies the maximum hosts/flows ntopng should expect to have on each interface

Configure ntopng/nProbe

- All the set of available options can be found in the ntopng/nprobe documentations
- https://www.ntop.org/guides/nprobe/cli_options.html
- https://www.ntop.org/guides/ntopng/cli_options/cli_options.html

Start ntopng/nProbe

- ntopng and nProbe can be start as services (suggested) by using systemctl or systemd (Linux)

- start/enable ntopng service:
  # sudo systemctl start ntopng
  # sudo systemctl enable ntopng

- start/enable nprobe service:
  # sudo systemctl start nprobe
  # sudo systemctl enable nprobe

- After starting both services connect to the ntopng Web Interface

  http(s)://NTOPNG_HOST_IP:PORT

- example (the default ntopng port is 3000):

  http://192.168.2.97:3000/

Start ntopng/nProbe

- In case multiple nProbes needs to be started, one config. file per nprobe needs to be created (in /etc/nprobe/ folder)
- Each config. file needs to have a different name, always starting with 'nprobe-', e.g.:
  - nprobe-eno1.conf
  - nprobe-router1.conf
- For each config, start a different service:
  # sudo systemctl start nprobe@eno1
  # sudo systemctl enable nprobe@eno1

  # sudo systemctl start nprobe@router1
  # sudo systemctl enable nprobe@router1

Result

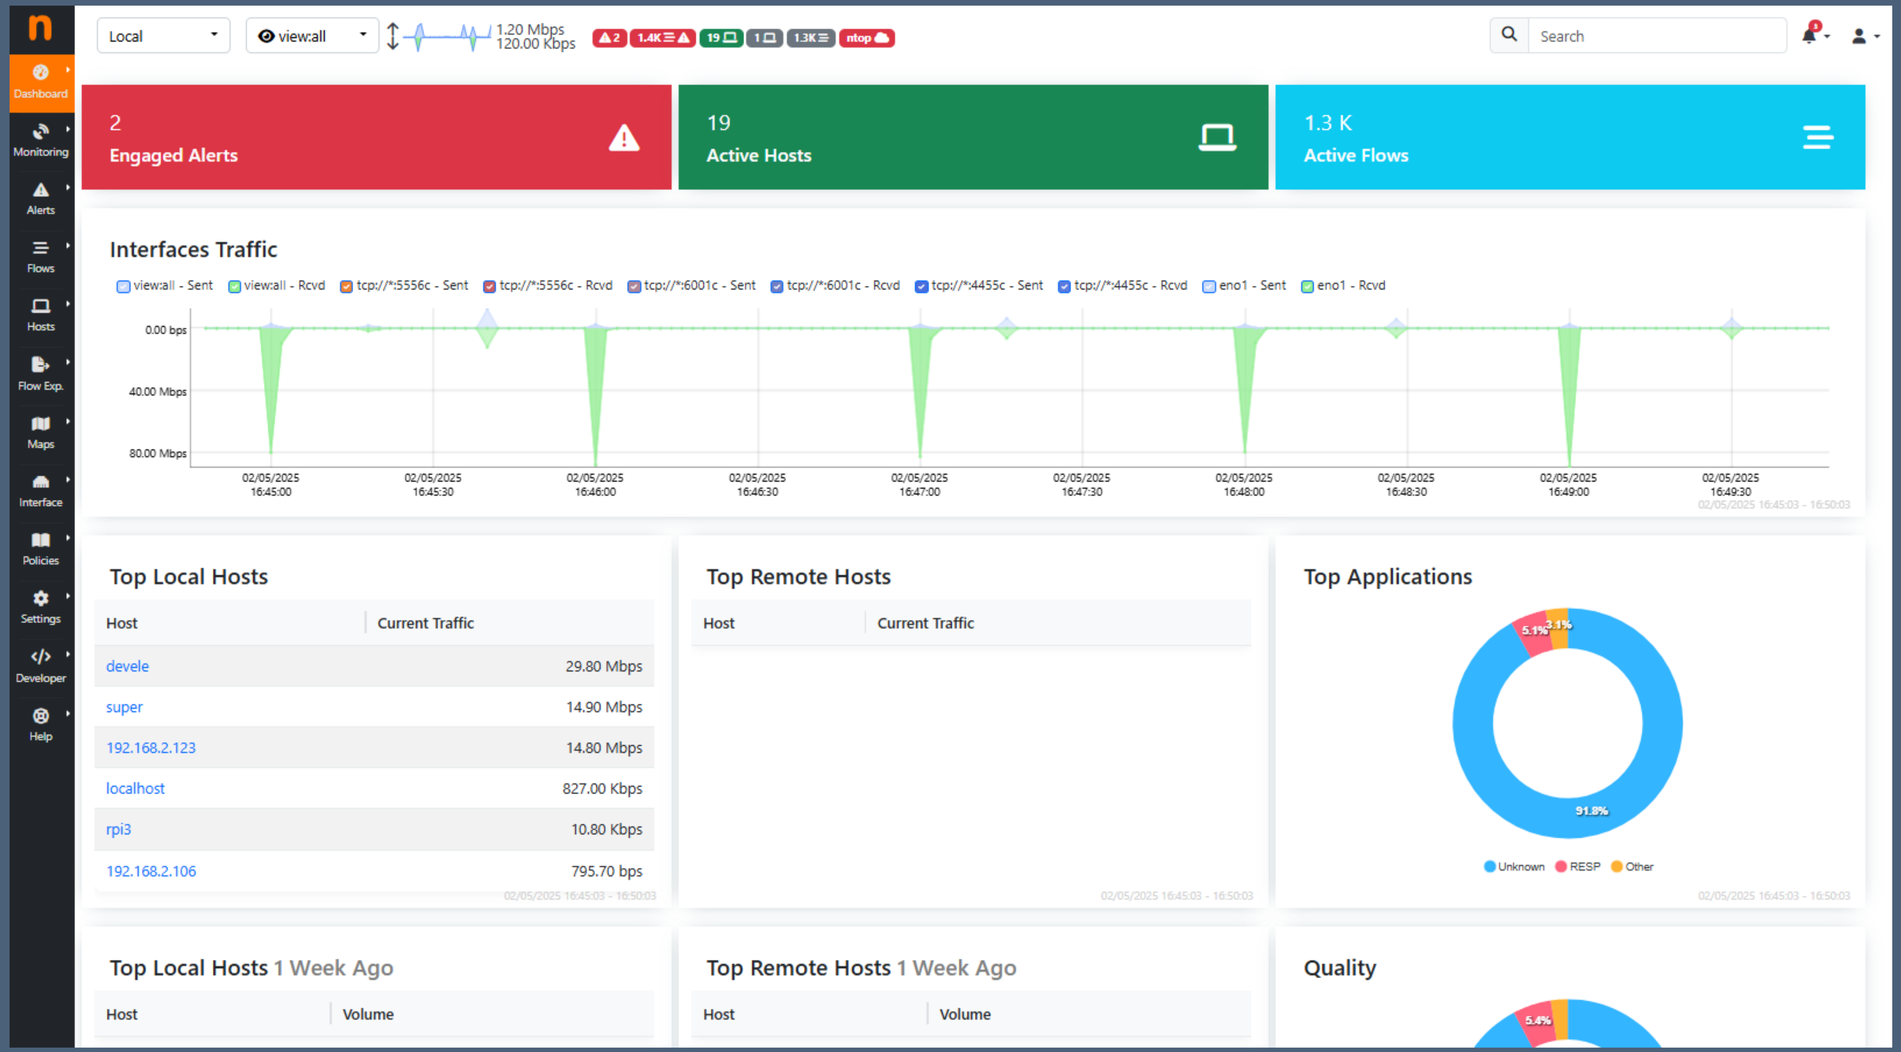
Local
view:all
1.20 Mbps
120.00 Kbps
Search
Dashboard
Monitoring
Alerts
Flows
Hosts
Flow Exp.
Maps
Interface
Policies
Settings
Developer
Help
2
Engaged Alerts
19
Active Hosts
1.3 K
Active Flows
Interfaces Traffic
view:all - Sent
view:all - Rcvd
tcp://*:5556c - Sent
tcp://*:5556c - Rcvd
tcp://*:6001c - Sent
tcp://*:6001c - Rcvd
tcp://*:4455c - Sent
tcp://*:4455c - Rcvd
eno1 - Sent
eno1 - Rcvd
0.00 bps
40.00 Mbps
80.00 Mbps
02/05/2025 16:45:00
02/05/2025 16:45:30
02/05/2025 16:46:00
02/05/2025 16:46:30
02/05/2025 16:47:00
02/05/2025 16:47:30
02/05/2025 16:48:00
02/05/2025 16:48:30
02/05/2025 16:49:00
02/05/2025 16:49:30
02/05/2025 16:45:03 - 16:50:03
Top Local Hosts
Host
Current Traffic
develе
29.80 Mbps
super
14.90 Mbps
192.168.2.123
14.80 Mbps
localhost
827.00 Kbps
rpi3
10.80 Kbps
192.168.2.106
795.70 bps
02/05/2025 16:45:03 - 16:50:03
Top Remote Hosts
Host
Current Traffic
02/05/2025 16:45:03 - 16:50:03
Top Applications
5.1%
3.1%
91.8%
Unknown
RESP
Other
02/05/2025 16:45:03 - 16:50:03
Top Local Hosts 1 Week Ago
Host
Volume
Top Remote Hosts 1 Week Ago
Host
Volume
Quality
5.4%

More Tuning: ClickHouse

- Now it is possible to see all the traffic, however we can do much more!
- Is it required to have those flows forwarded by the devices (switches/routers) accessible even in the future (for example, be able to access data from a week ago)?
- Configure the -F option in the ntopng configuration file
- ntopng is able to export flows to various DB:
  - ClickHouse (suggested)
  - Syslog
  - Kafka
  - Elasticsearch

ClickHouse

- The option is:
  - clickhouse;<host[@[<tcp-port>,]<mysql-port]|socket>;<dbname>;<user>;<pw>
- Add to the configuration file the right option (in case of localhost as host, standard port and standard username/password, simply add 'clickhouse'):
  - -F=clickhouse

Historical Flows

- Now in the Flows menu, a new option should be available: Historical
- Here it is possible to navigate through already expired flows

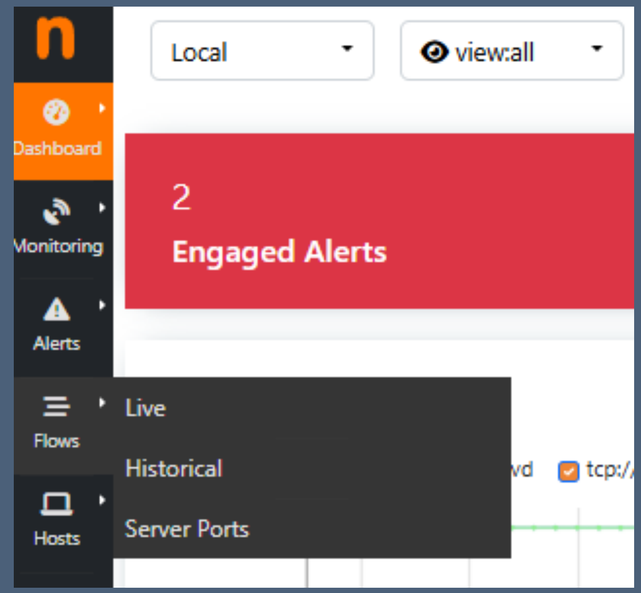

Historical Flows

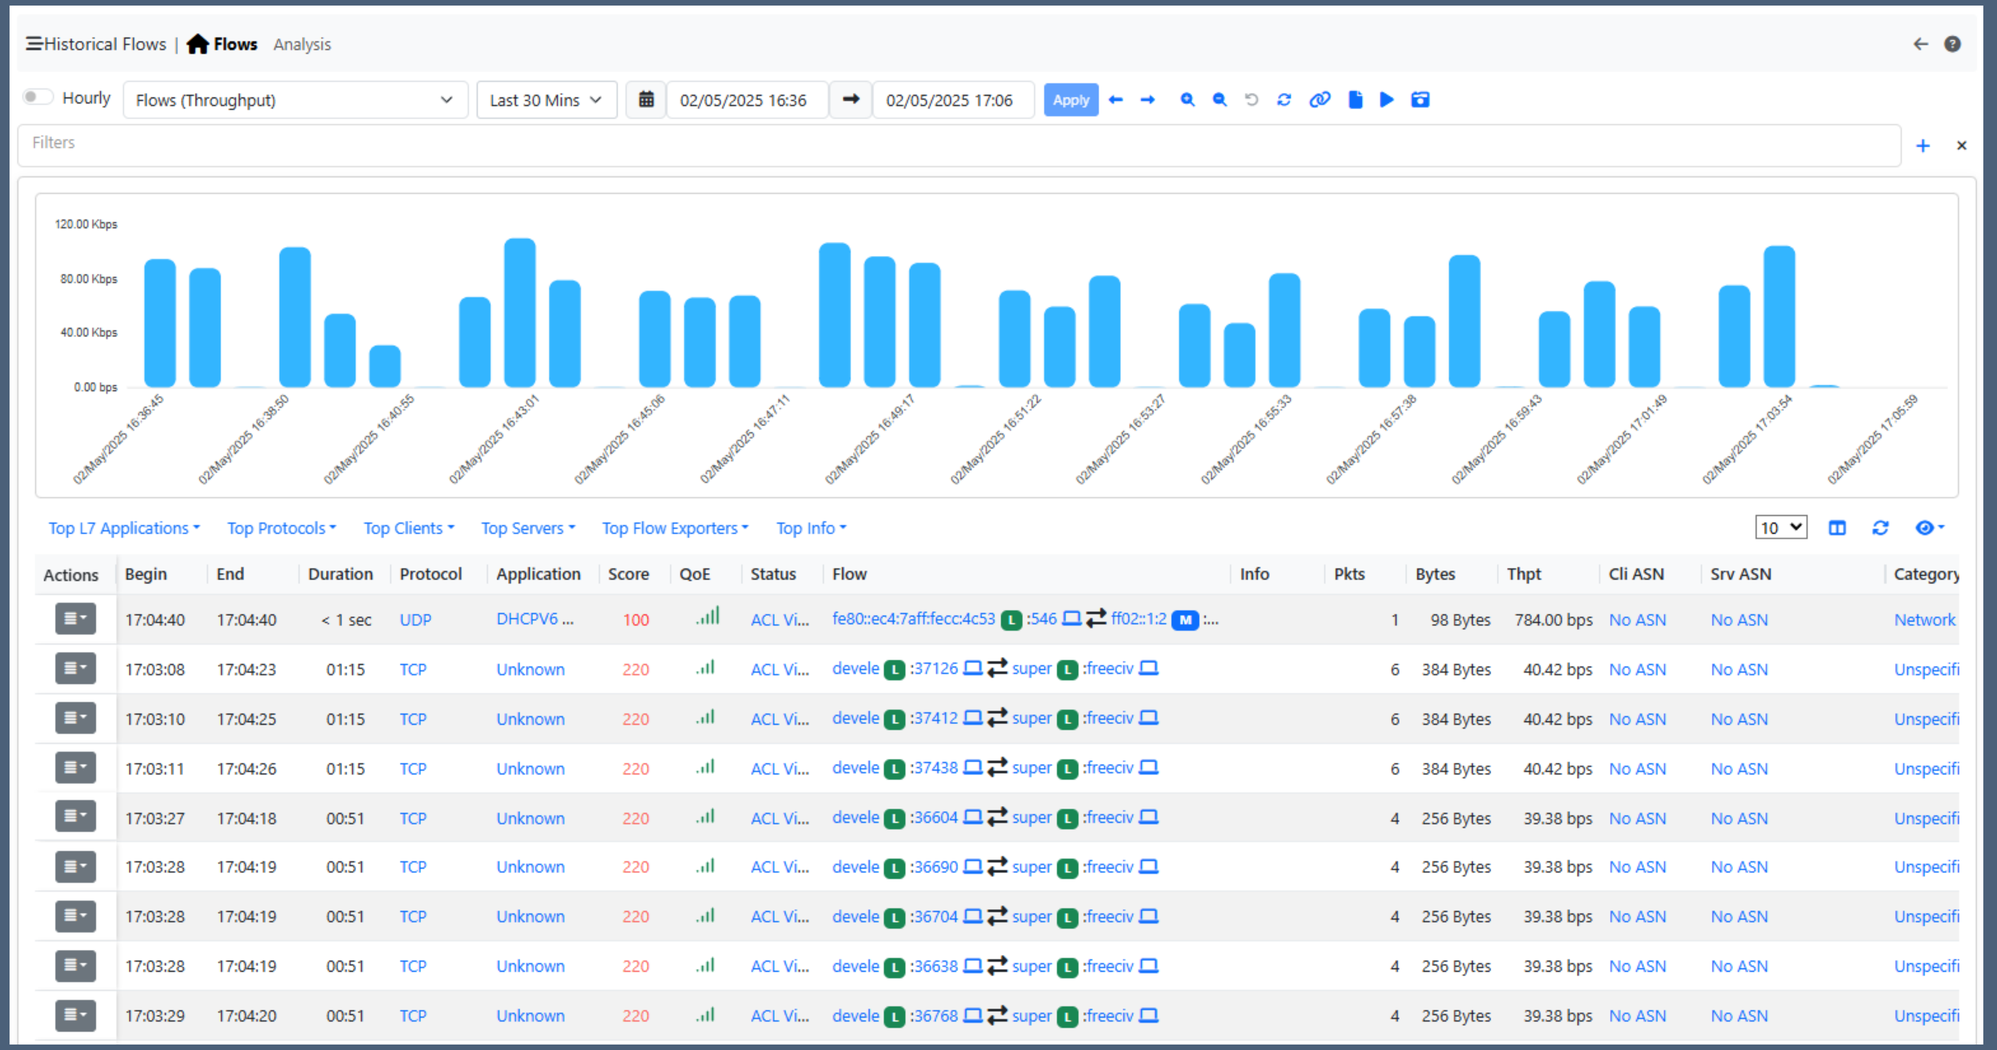

Actions	Begin	End	Duration	Protocol	Application	Score	QoE	Status	Flow	Info	Pkts	Bytes	Thpt	Cli ASN	Srv ASN	Category
	17:04:40	17:04:40	< 1 sec	UDP	DHCPV6 ...	100		ACL Vi...	fe80::ec4:7aff:fecc:4c53 L :546 ff02::1:2 M :...		1	98 Bytes	784.00 bps	No ASN	No ASN	Network
	17:03:08	17:04:23	01:15	TCP	Unknown	220		ACL Vi...	develo L :37126 super L :freeciv		6	384 Bytes	40.42 bps	No ASN	No ASN	Unspecifi
	17:03:10	17:04:25	01:15	TCP	Unknown	220		ACL Vi...	develo L :37412 super L :freeciv		6	384 Bytes	40.42 bps	No ASN	No ASN	Unspecifi
	17:03:11	17:04:26	01:15	TCP	Unknown	220		ACL Vi...	develo L :37438 super L :freeciv		6	384 Bytes	40.42 bps	No ASN	No ASN	Unspecifi
	17:03:27	17:04:18	00:51	TCP	Unknown	220		ACL Vi...	develo L :36604 super L :freeciv		4	256 Bytes	39.38 bps	No ASN	No ASN	Unspecifi
	17:03:28	17:04:19	00:51	TCP	Unknown	220		ACL Vi...	develo L :36690 super L :freeciv		4	256 Bytes	39.38 bps	No ASN	No ASN	Unspecifi
	17:03:28	17:04:19	00:51	TCP	Unknown	220		ACL Vi...	develo L :36704 super L :freeciv		4	256 Bytes	39.38 bps	No ASN	No ASN	Unspecifi
	17:03:28	17:04:19	00:51	TCP	Unknown	220		ACL Vi...	develo L :36638 super L :freeciv		4	256 Bytes	39.38 bps	No ASN	No ASN	Unspecifi
	17:03:29	17:04:20	00:51	TCP	Unknown	220		ACL Vi...	develo L :36768 super L :freeciv		4	256 Bytes	39.38 bps	No ASN	No ASN	Unspecifi

Historical Flows

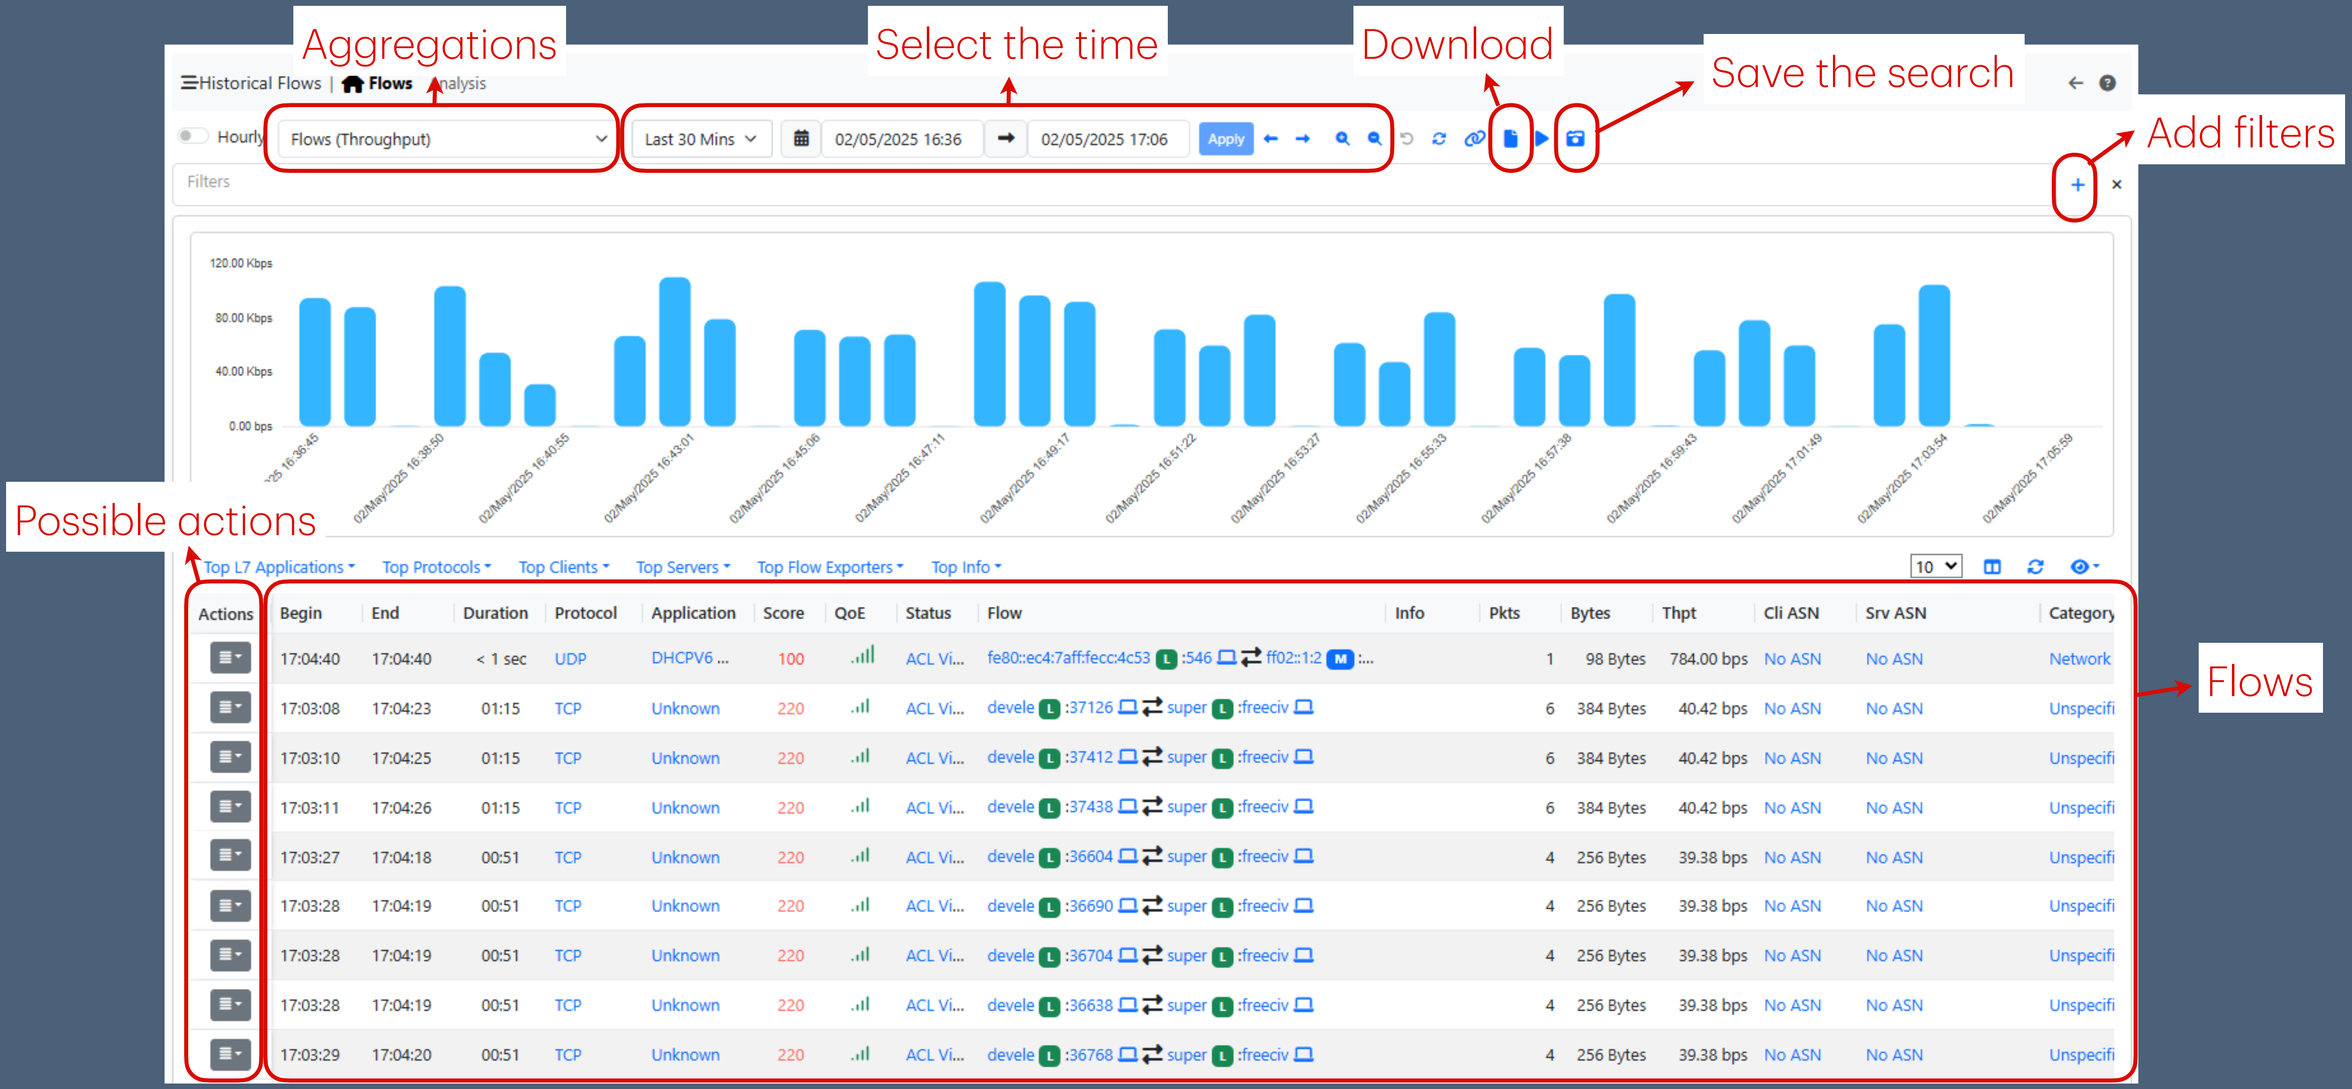

Tuning ClickHouse

- It is important to know that each flow, when saved on the DB, occupy some disk, so it is important to tune the Data Retention

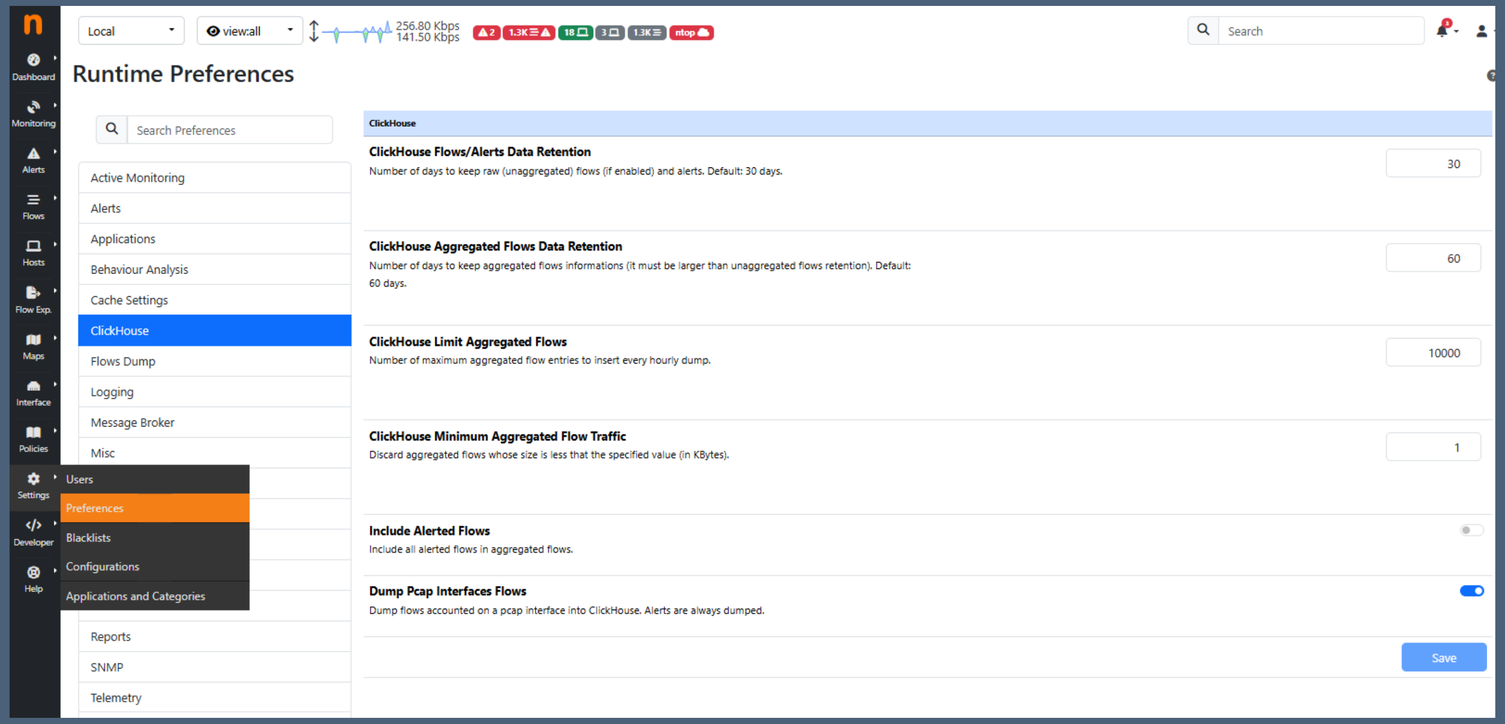

Tuning ClickHouse

- Also sometimes ClickHouse uses a lot of RAM/CPU/Disk without reason, so it's highly suggested to set a couple of things

- Check the Disk used by ClickHouse:
  # clickhouse-client
  # use ntopng
  # SELECT table, formatReadableSize(sum(bytes)) AS size, min(min_date) AS min_date, max(max_date) AS max_date FROM system.parts WHERE active GROUP BY table

- Reduce the table TTL of the system tables:
  # ALTER TABLE system.XXX MODIFY TTL event_date + INTERVAL 3 DAY;
  # ALTER TABLE system.query_log MODIFY TTL event_date + INTERVAL 3 DAY;
  # ALTER TABLE system.asynchronous_metric_log MODIFY TTL event_date + INTERVAL 3 DAY;
  # ALTER TABLE system.metric_log MODIFY TTL event_date + INTERVAL 3 DAY;
  # ALTER TABLE system.trace_log MODIFY TTL event_date + INTERVAL 3 DAY;

- https://github.com/ntop/ntopng/blob/dev/doc/README.clickhouse.md

Ready

- Everything is done now!
- Explore the currently active flows
- Explore the historical flows
- Explore the local/remote hosts

More Analysis

- An important step in the network analysis are also timeseries
- There are a lot of timeseries in ntopng and they can greatly enchant the network analysis

Timeseries

- Before jumping to the timeseries, start collecting the interesting ones
- Jump to the preferences and enable/disable the important ones

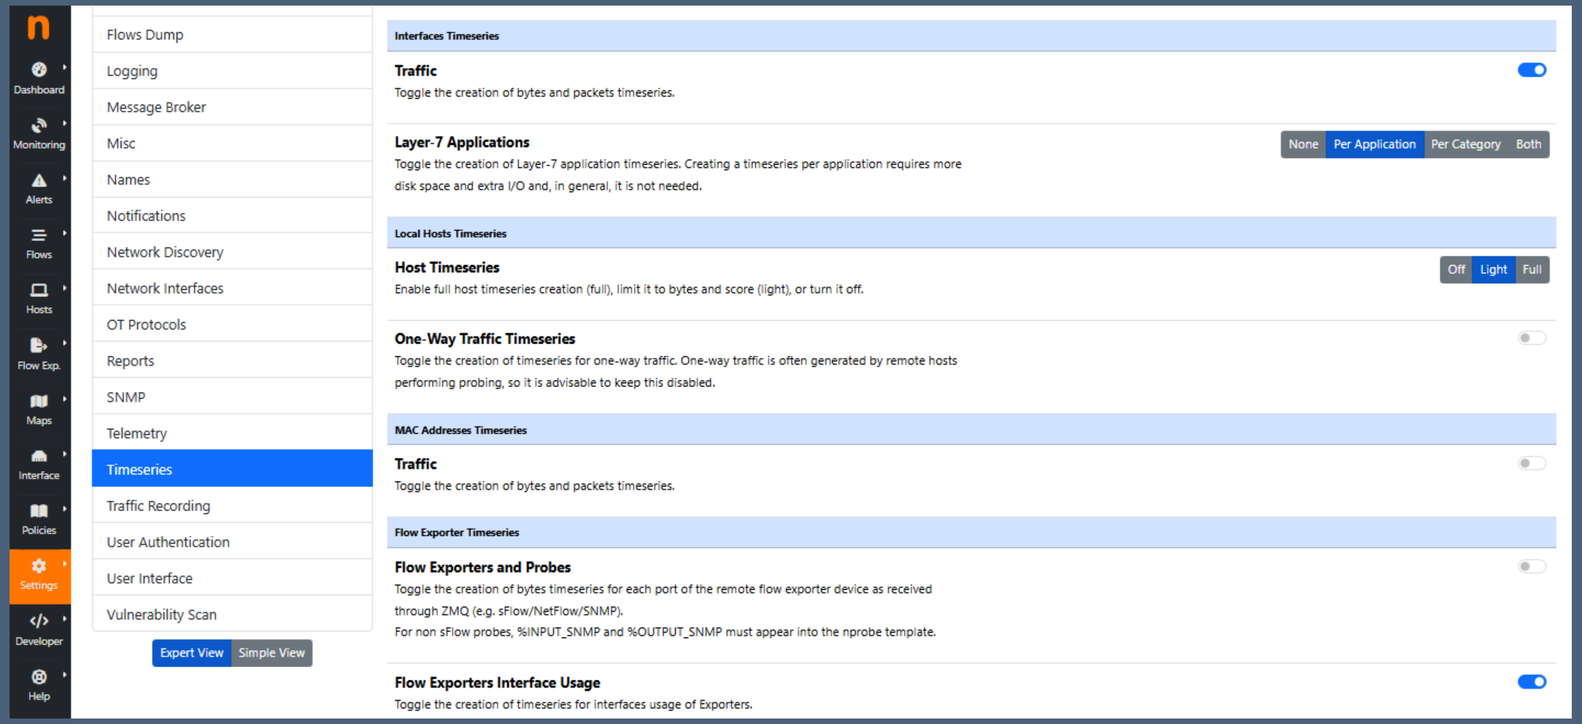

Timeseries Drivers

- Two drivers are available in ntopng to store timeseries:
  - RRD
  - InfluxDB
- Different use cases:
  - RRD: faster and uses very few RAM/CPU/Disk, however it can only be on the same machine where ntopng is
  - InfluxDB: slower and in can use a lot of resources, however in can be placed on a different machine; also it is a lot more flexible

Timeseries

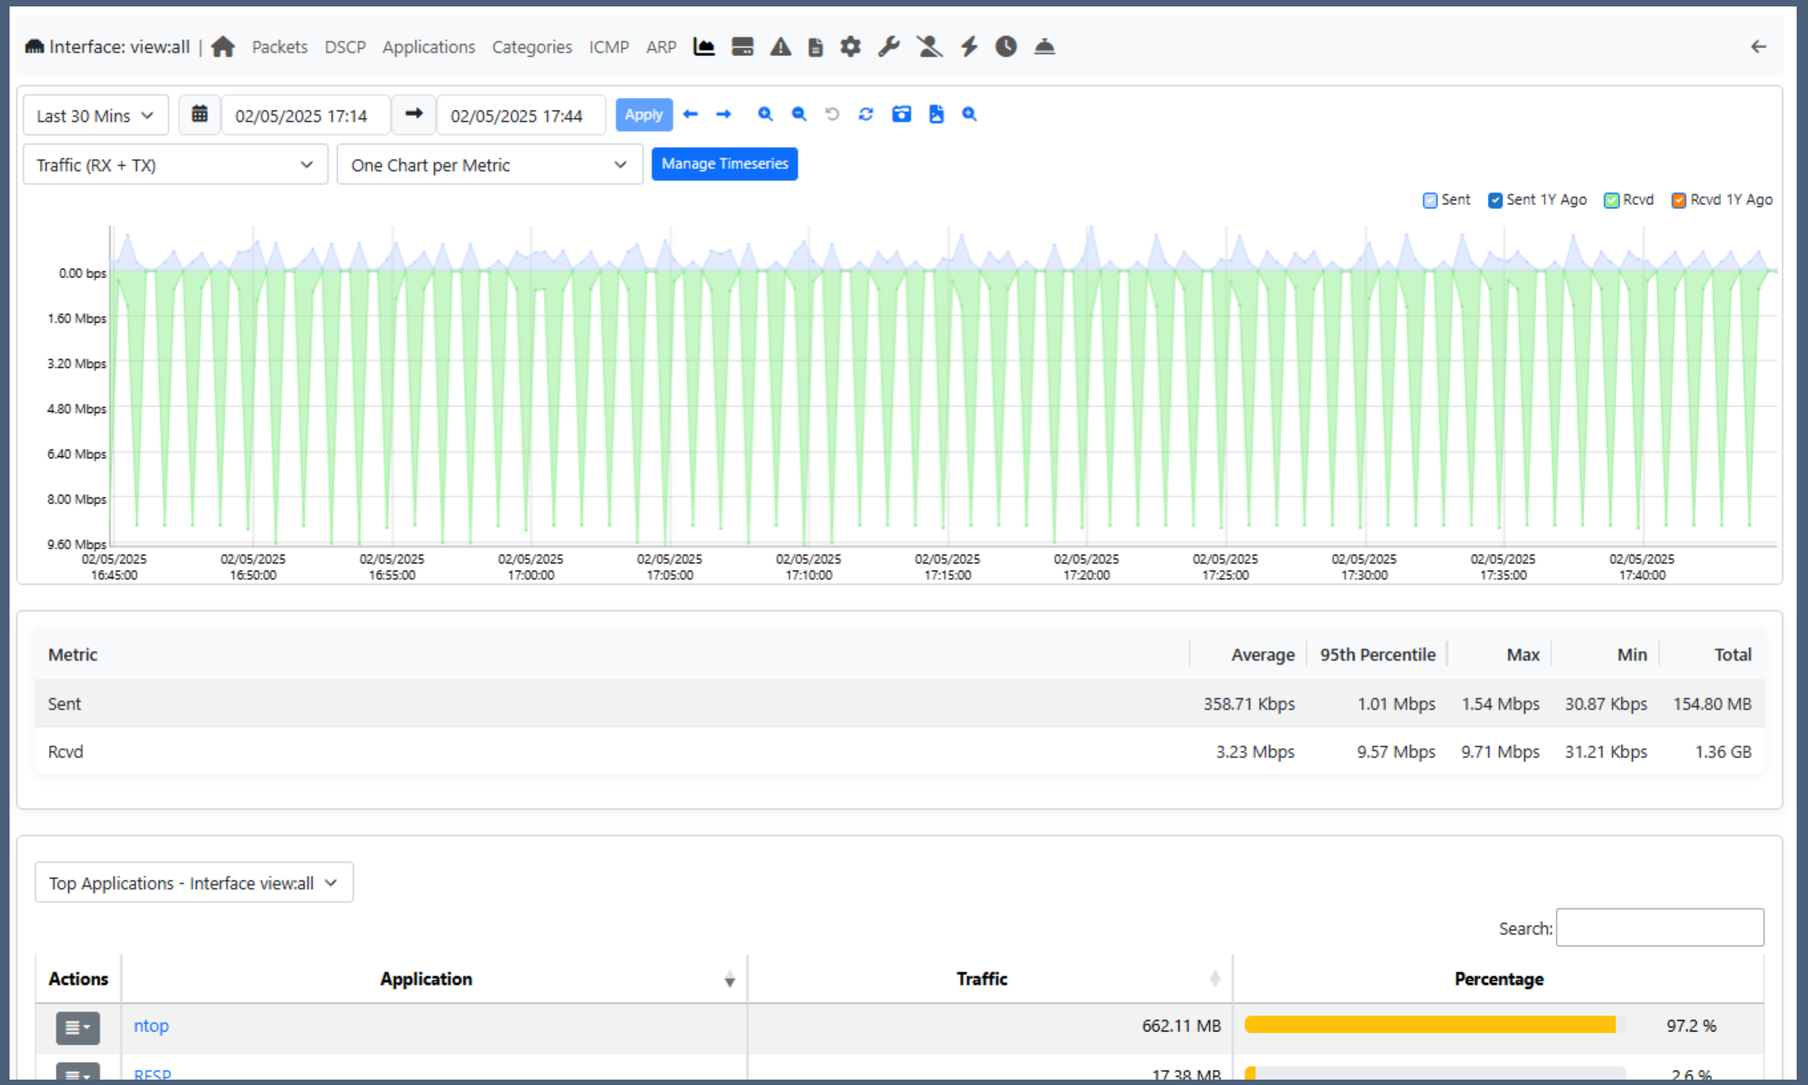
Interface: view:all | Packets DSCP Applications Categories ICMP ARP
Last 30 Mins
02/05/2025 17:14
02/05/2025 17:44
Apply
Traffic (RX + TX)
One Chart per Metric
Manage Timeseries
Sent
Sent 1Y Ago
Rcvd
Rcvd 1Y Ago
0.00 bps
1.60 Mbps
3.20 Mbps
4.80 Mbps
6.40 Mbps
8.00 Mbps
9.60 Mbps
02/05/2025 16:45:00
02/05/2025 16:50:00
02/05/2025 16:55:00
02/05/2025 17:00:00
02/05/2025 17:05:00
02/05/2025 17:10:00
02/05/2025 17:15:00
02/05/2025 17:20:00
02/05/2025 17:25:00
02/05/2025 17:30:00
02/05/2025 17:35:00
02/05/2025 17:40:00
Metric | Average | 95th Percentile | Max | Min | Total
Sent | 358.71 Kbps | 1.01 Mbps | 1.54 Mbps | 30.87 Kbps | 154.80 MB
Rcvd | 3.23 Mbps | 9.57 Mbps | 9.71 Mbps | 31.21 Kbps | 1.36 GB
Top Applications - Interface view:all
Search:
Actions | Application | Traffic | Percentage
ntop | 662.11 MB | 97.2 %
RESP | 17.38 MB | 2.6 %

Timeseries

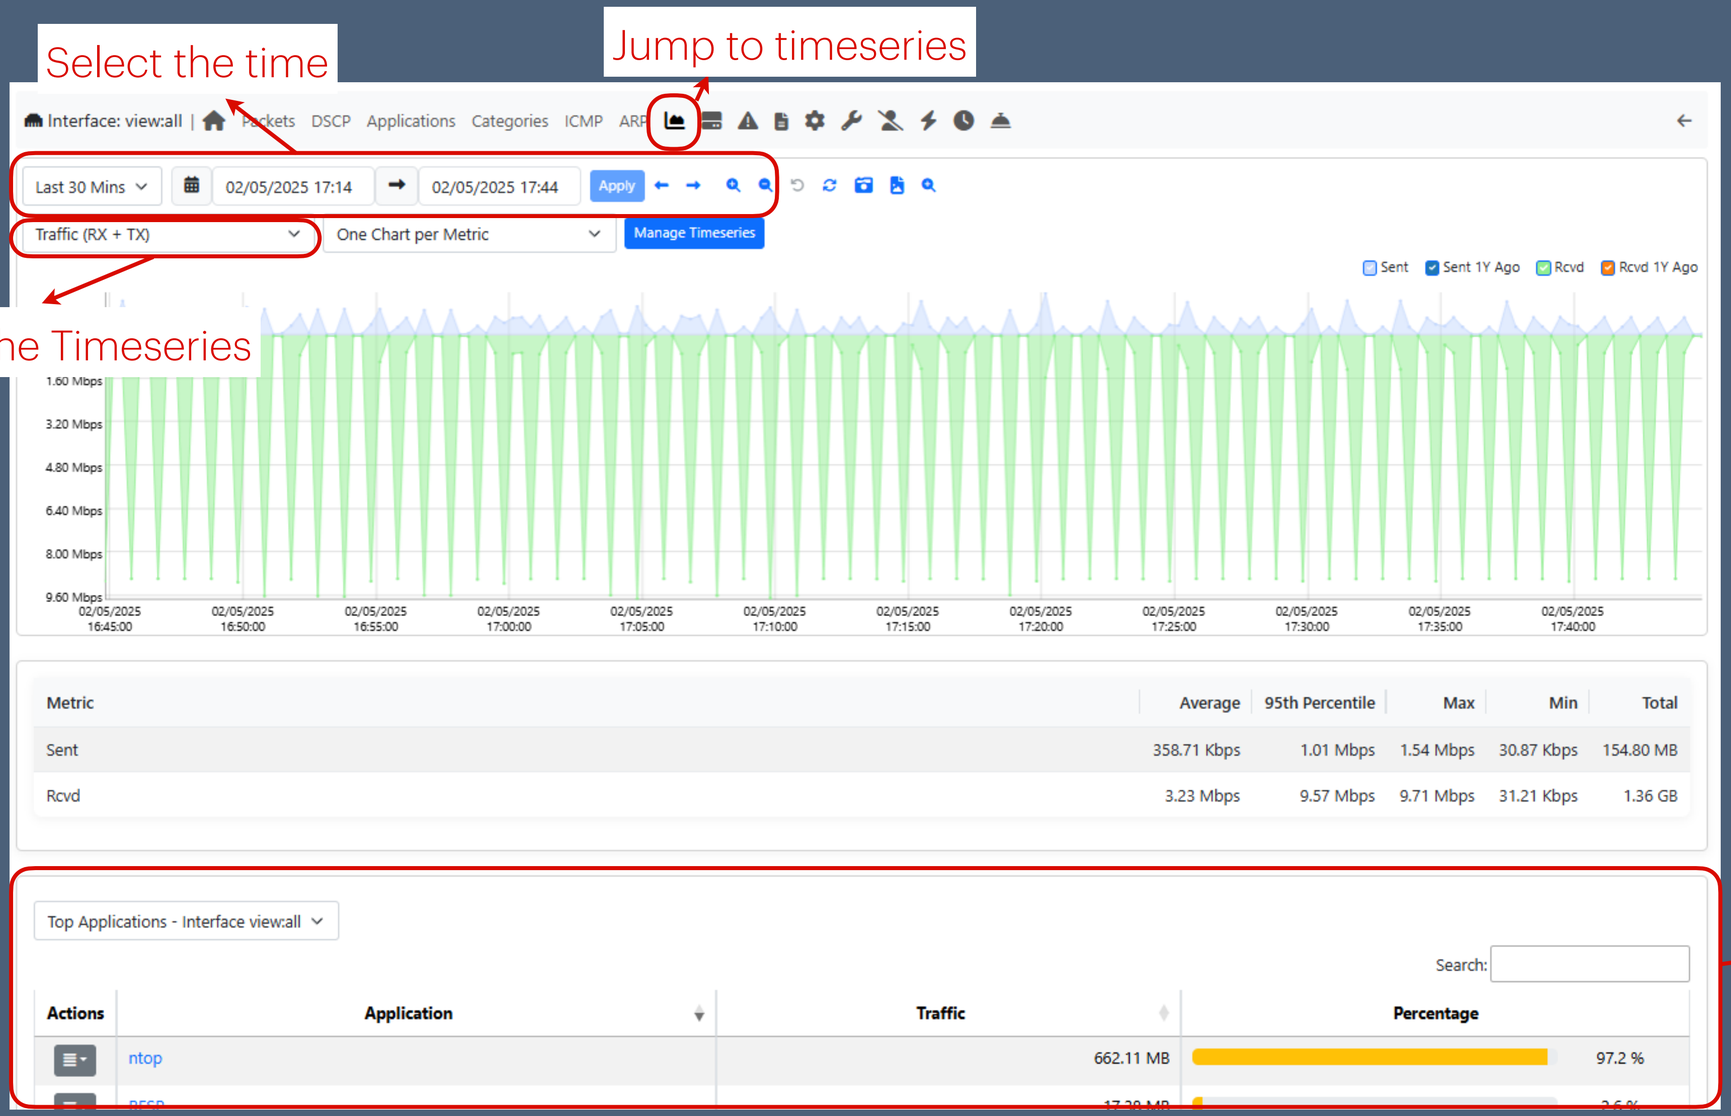
Select the time
Jump to timeseries
Select the Timeseries
Interface: view:all
Packets
DSCP
Applications
Categories
ICMP
ARP
Last 30 Mins
02/05/2025 17:14
02/05/2025 17:44
Apply
Traffic (RX + TX)
One Chart per Metric
Manage Timeseries
Sent
Sent 1Y Ago
Rcvd
Rcvd 1Y Ago
1.60 Mbps
3.20 Mbps
4.80 Mbps
6.40 Mbps
8.00 Mbps
9.60 Mbps
02/05/2025 16:45:00
02/05/2025 16:50:00
02/05/2025 16:55:00
02/05/2025 17:00:00
02/05/2025 17:05:00
02/05/2025 17:10:00
02/05/2025 17:15:00
02/05/2025 17:20:00
02/05/2025 17:25:00
02/05/2025 17:30:00
02/05/2025 17:35:00
02/05/2025 17:40:00
Metric
Average
95th Percentile
Max
Min
Total
Sent
358.71 Kbps
1.01 Mbps
1.54 Mbps
30.87 Kbps
154.80 MB
Rcvd
3.23 Mbps
9.57 Mbps
9.71 Mbps
31.21 Kbps
1.36 GB
Top Applications - Interface view:all
Search:
Actions
Application
Traffic
Percentage
ntop
662.11 MB
97.2 %
Tops (if available)

Enchant the analysis

- It is possible to further increase the analysis and 'control' over the network
- Why only check data/flows?
- Move to also check the devices, by using SNMP

SNMP

- It is possible to monitor the status of the various devices by using SNMP
- ntopng supports various MIBs:
  - MIB II
  - LLDP/CDP MIB
  - Bridge MIB
  - netSNMP
  - Cisco QoS
  - Cisco CPU/Memory MIB (in development)

Advantages of using SNMP

- ntopng polls the bridge MIB that is used to discover the MAC addresses observed on a network interface
- ntopng allows interface traffic to be compared for similarity; if one interface is miss behaving and has a similarity with an other one, are both interfaces affected?
- With the support of LLDP and CDP it's possible to build maps of devices

Advantages of using SNMP

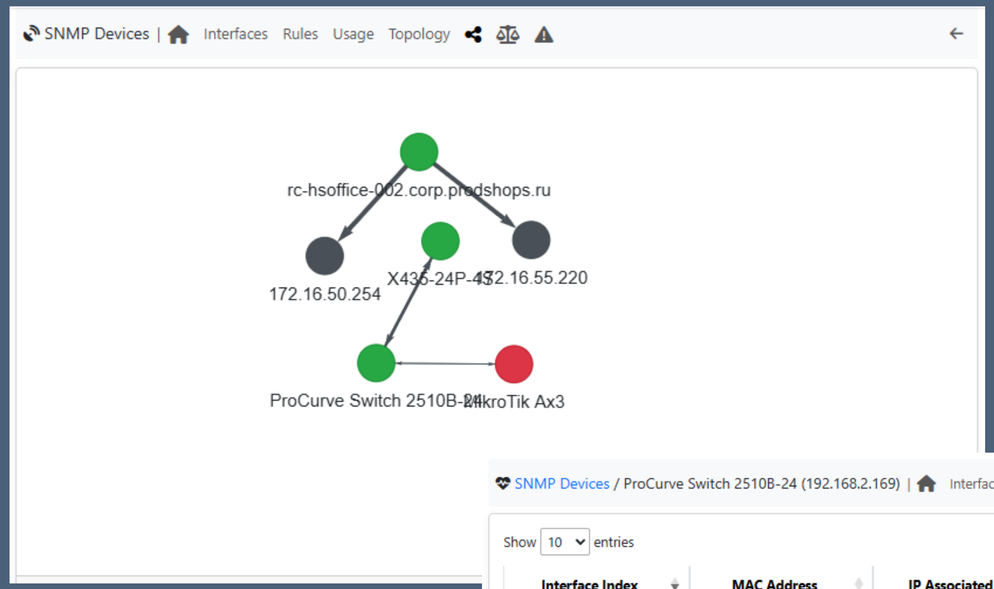

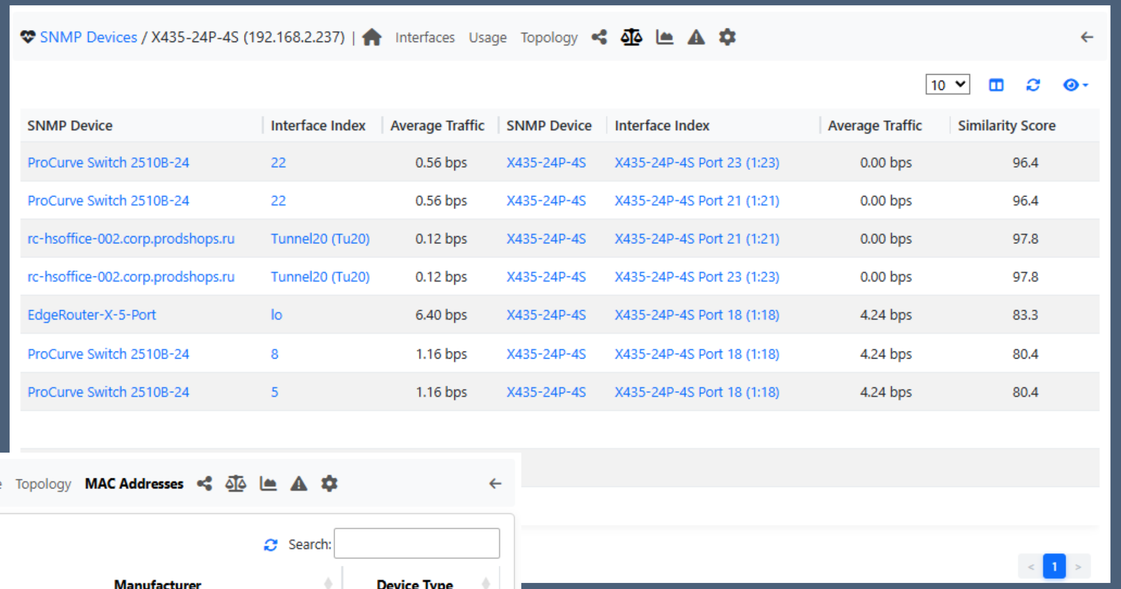

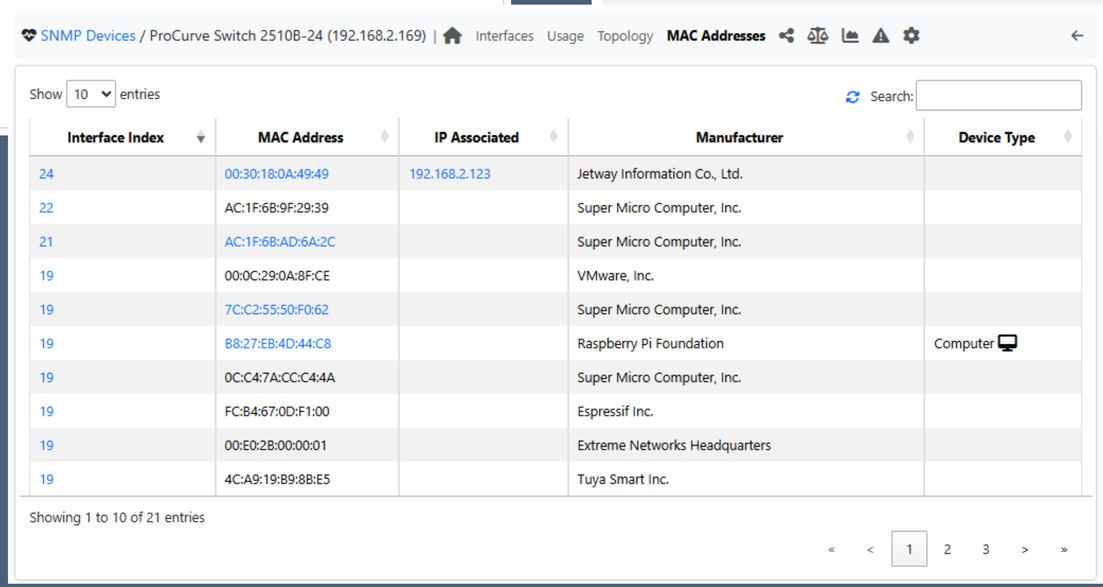

… and more SNMP

- Also in the latest version it is possible, by using SNMP, to poll the speed of the various SNMP interfaces
- Comparing the speeds with the bytes sent/rcvd collected from the other MIBs, it is possible to understand if an interface is congested

SNMP Usage

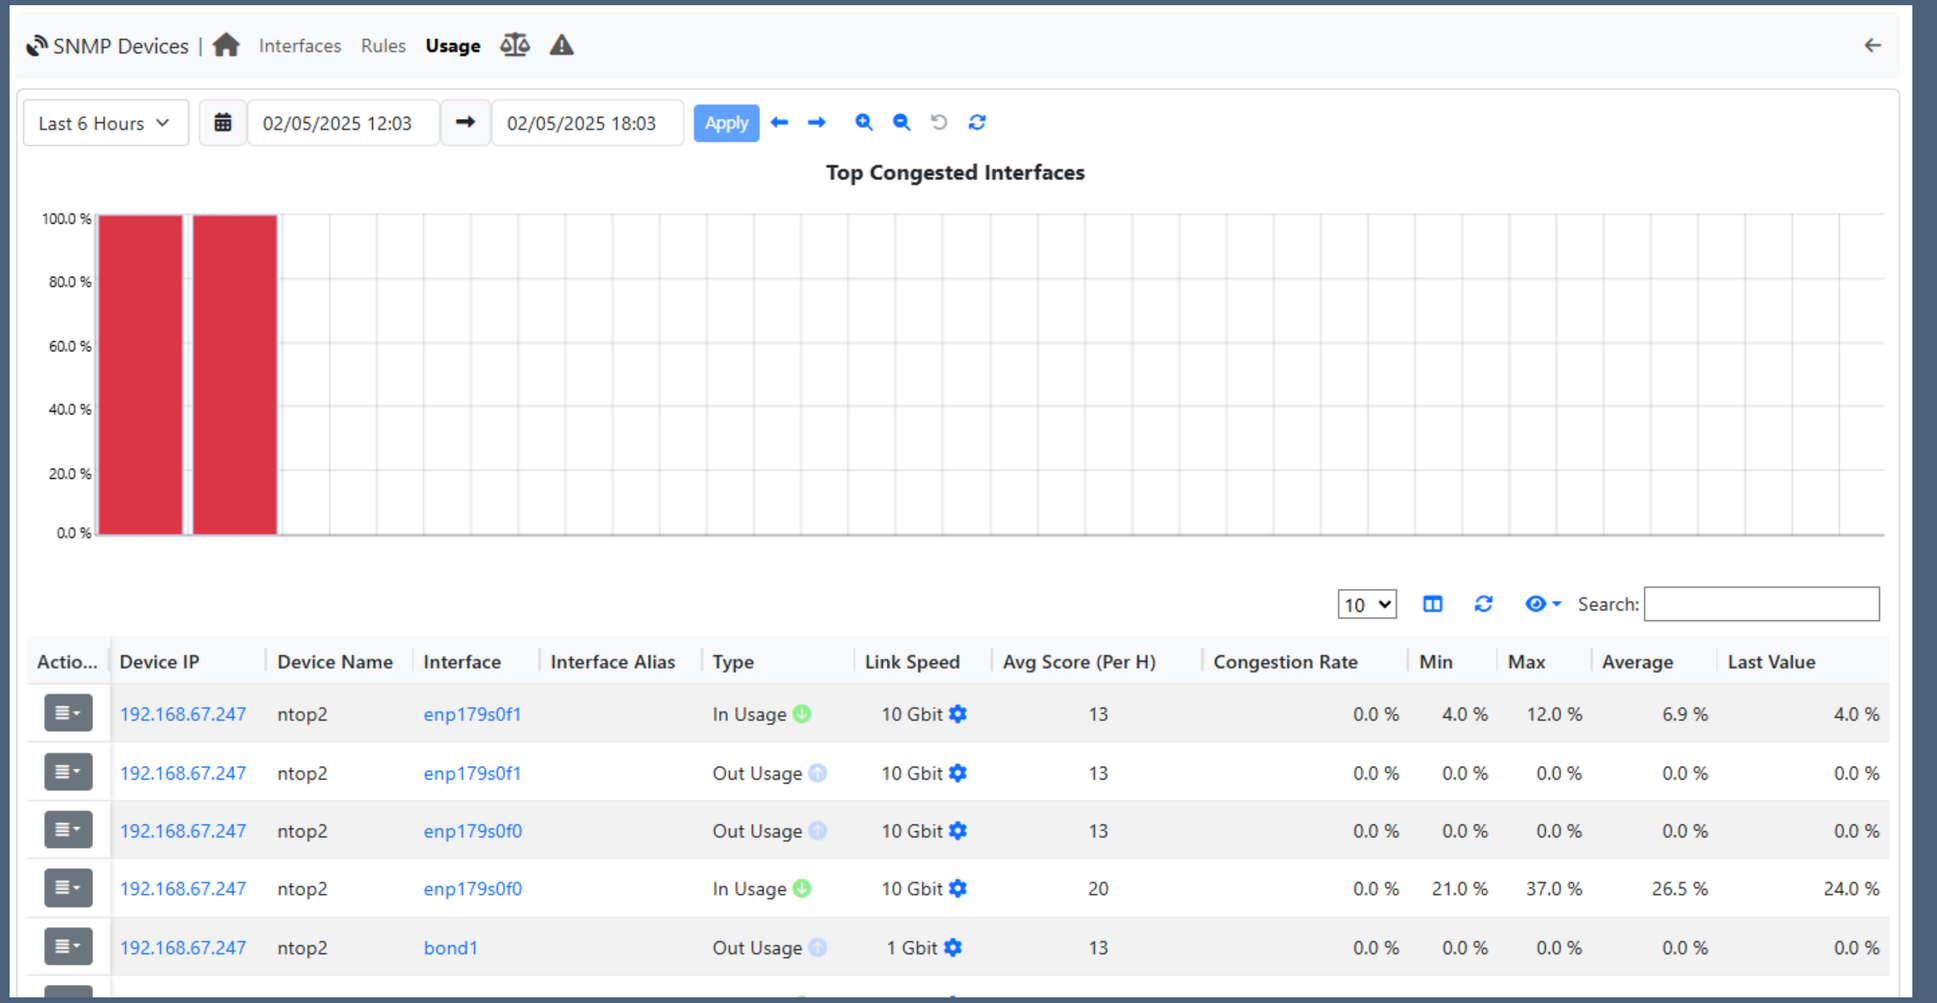

Actio...	Device IP	Device Name	Interface	Interface Alias	Type	Link Speed	Avg Score (Per H)	Congestion Rate	Min	Max	Average	Last Value
	192.168.67.247	ntop2	enp179s0f1		In Usage	10 Gbit	13	0.0 %	4.0 %	12.0 %	6.9 %	4.0 %
	192.168.67.247	ntop2	enp179s0f1		Out Usage	10 Gbit	13	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %
	192.168.67.247	ntop2	enp179s0f0		Out Usage	10 Gbit	13	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %
	192.168.67.247	ntop2	enp179s0f0		In Usage	10 Gbit	20	0.0 %	21.0 %	37.0 %	26.5 %	24.0 %
	192.168.67.247	ntop2	bond1		Out Usage	1 Gbit	13	0.0 %	0.0 %	0.0 %	0.0 %	0.0 %

SNMP Usage

- If the interface speed is collected correctly by SNMP (it is also possible to manually configure it), ntopng is able to detect congested interfaces
- An interface is marked as congested if the usage (Link Speed / Bytes sent-rcvd) is higher or equal than 76%
- Also think about the previous configurations done...
- YES! ClickHouse!
- By comparing the congested interface time period and the historical data, it is possible (if correctly configured) who congested the link

SNMP Usage

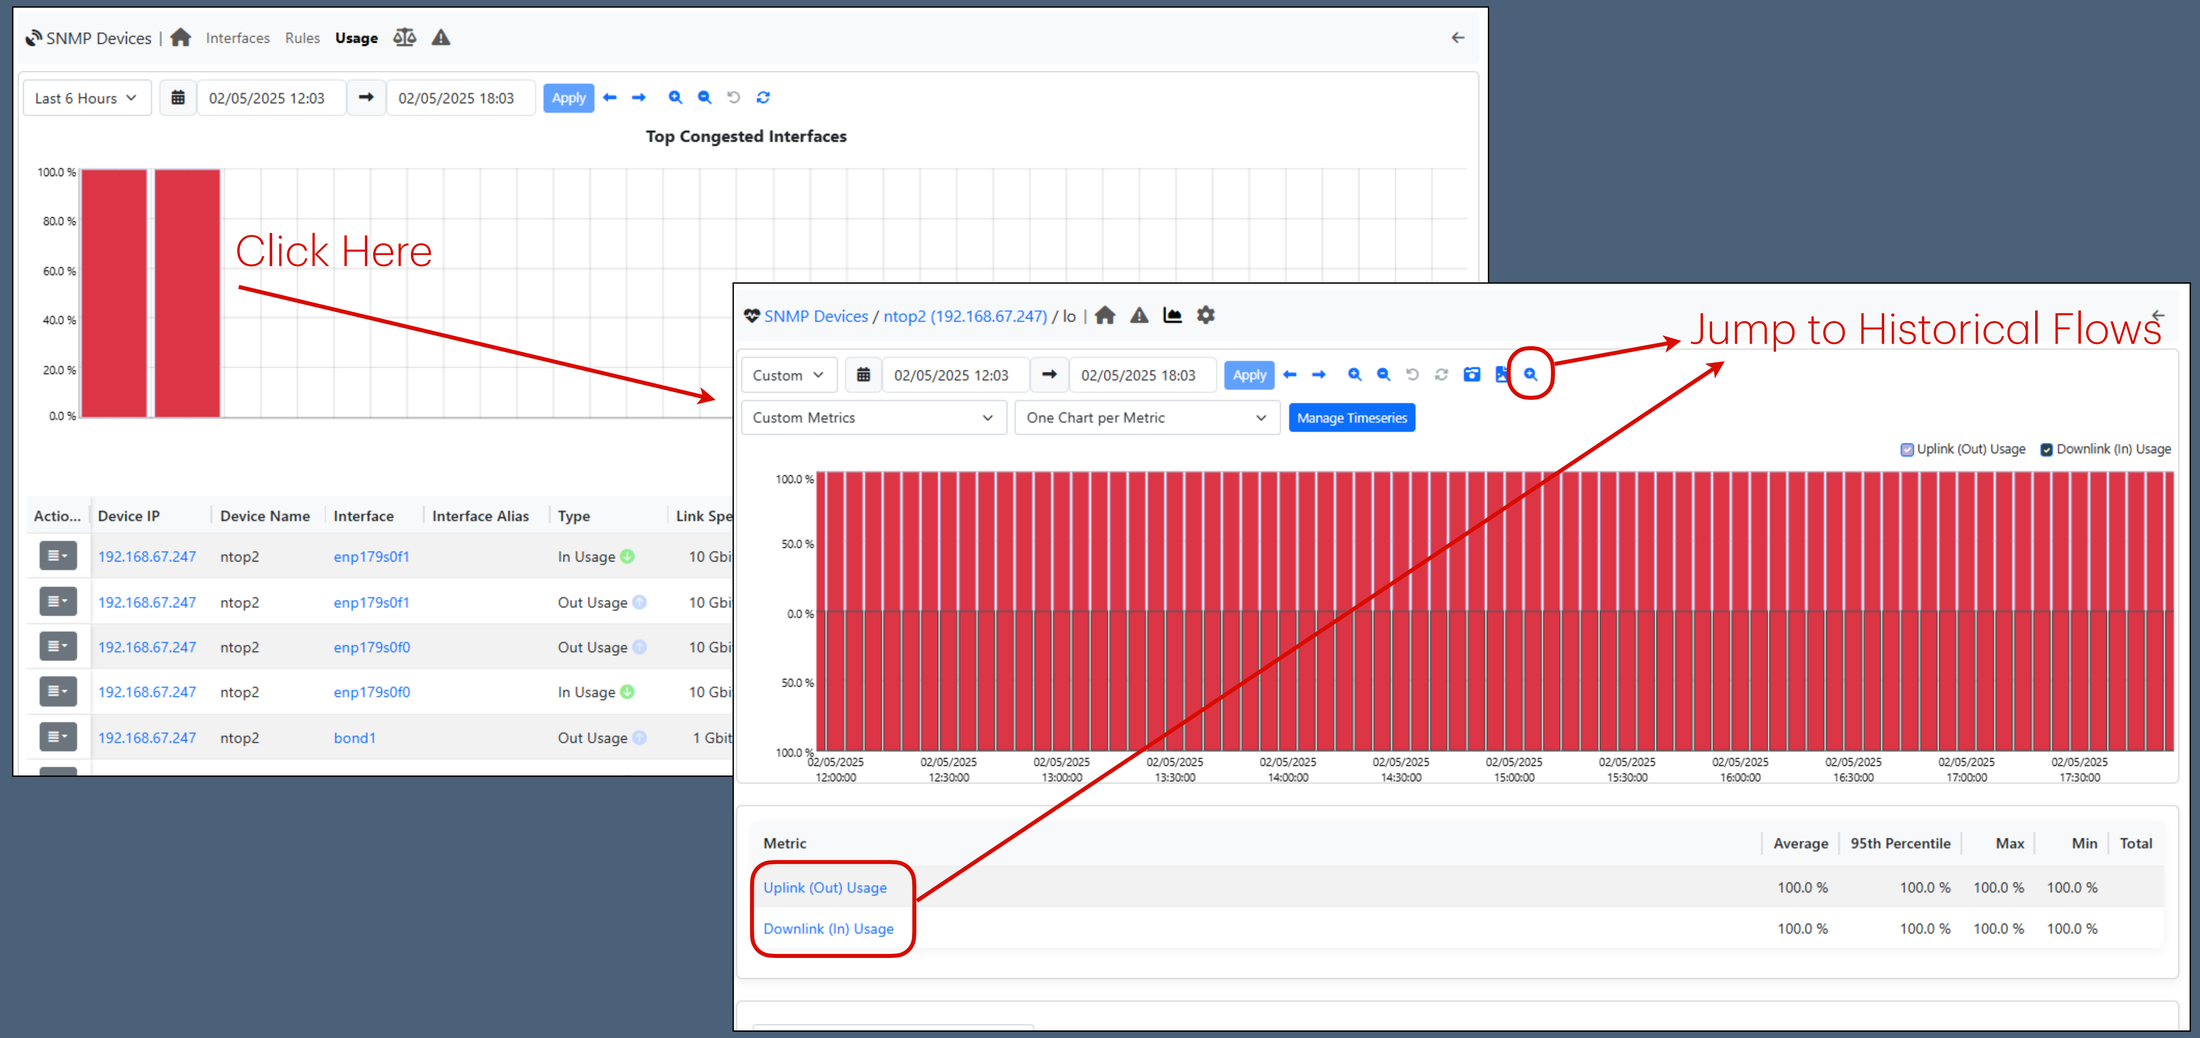
SNMP Devices | Interfaces Rules Usage
Last 6 Hours
02/05/2025 12:03
02/05/2025 18:03
Apply
Top Congested Interfaces
100.0 %
80.0 %
60.0 %
40.0 %
20.0 %
0.0 %
Click Here
Actio... Device IP Device Name Interface Interface Alias Type Link Spe
192.168.67.247 ntop2 enp179s0f1 In Usage 10 Gbi
192.168.67.247 ntop2 enp179s0f1 Out Usage 10 Gbi
192.168.67.247 ntop2 enp179s0f0 Out Usage 10 Gbi
192.168.67.247 ntop2 enp179s0f0 In Usage 10 Gbi
192.168.67.247 ntop2 bond1 Out Usage 1 Gbit
SNMP Devices / ntop2 (192.168.67.247) / lo
Jump to Historical Flows
Custom
02/05/2025 12:03
02/05/2025 18:03
Apply
Custom Metrics
One Chart per Metric
Manage Timeseries
Uplink (Out) Usage
Downlink (In) Usage
100.0 %
50.0 %
0.0 %
50.0 %
100.0 %
02/05/2025 12:00:00
02/05/2025 12:30:00
02/05/2025 13:00:00
02/05/2025 13:30:00
02/05/2025 14:00:00
02/05/2025 14:30:00
02/05/2025 15:00:00
02/05/2025 15:30:00
02/05/2025 16:00:00
02/05/2025 16:30:00
02/05/2025 17:00:00
02/05/2025 17:30:00
Metric Average 95th Percentile Max Min Total
Uplink (Out) Usage 100.0 % 100.0 % 100.0 % 100.0 %
Downlink (In) Usage 100.0 % 100.0 % 100.0 % 100.0 %

Conclusions

- We have seen how to properly configure nProbe/ntopng to analyze a network
- How to store the flows by using a DB (to be able to access those data in the future)
- How to use timeseries to have an idea of the traffic/activities done by hosts, interfaces, ...
- How to also support the devices in order to understand if there is someone filling the link

Questions

Thanks

https://github.com/ntop/ntopng

https://www.ntop.org/