nDPI and nProbe

Luca Deri <deri@ntop.org>
@lucaderi

Part I: nDPI

nDPI: A Recap

- nDPI is an open source toolkit that classifies traffic using DPI, deep packet inspection.

Layer 4 Protocol

Good or Bad?

TCP / HTTP

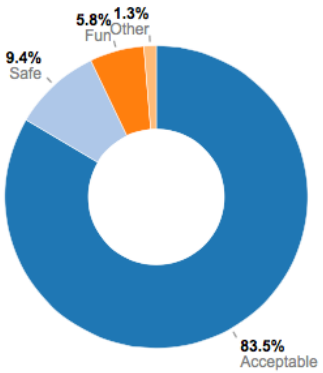

Layer 7 Protocol

https://github.com/ntop/nDPI

nDPI and Application Detection

```
nDPI supported protocols:
 Id Protocol               Layer_4    Nw_Proto Breed        Category
  0 Unknown                TCP        X        Unrated      Unspecified
  1 FTP_CONTROL            TCP        X        Unsafe       Download
  2 POP3                   TCP        X        Unsafe       Email
  3 SMTP                   TCP        X        Acceptable   Email
  4 IMAP                   TCP        X        Unsafe       Email
  5 DNS                    TCP/UDP    X        Acceptable   Network
…
291 MpegDash               TCP                 Acceptable   Media
292 Dazn                   TCP                 Fun          Streaming
293 GoTo                   TCP                 Acceptable   VoIP
294 RSH                    TCP        X        Unsafe       RemoteAccess
295 1kxun                  TCP                 Fun          Streaming
296 PGM                               X        Acceptable   Network
297 IP_PIM                            X        Acceptable   Network
298 collectd               UDP        X        Acceptable   System
```

nDPI and Traffic Analysis

Id	Risk	Severity	Score	CliScore	SrvScore
1	XSS Attack	Severe	250	225	25
2	SQL Injection	Severe	250	225	25
3	RCE Injection	Severe	250	225	25
4	Binary App Transfer	Severe	250	125	125
5	Known Proto on Non Std Port	Medium	50	25	25
6	Self-signed Cert	High	100	90	10
7	Obsolete TLS (v1.1 or older)	High	100	90	10
8	Weak TLS Cipher	High	100	90	10
9	TLS Cert Expired	High	100	10	90
10	TLS Cert Mismatch	High	100	50	50
11	HTTP Suspicious User-Agent	High	100	90	10
12	HTTP Numeric IP Address	Low	10	5	5
13	HTTP Suspicious URL	High	100	90	10
14	HTTP Suspicious Header	High	100	90	10
…					
39	Text With Non-Printable Chars	High	100	90	10
40	Possible Exploit	Severe	250	225	25
41	TLS Cert About To Expire	Medium	50	5	45
42	IDN Domain Name	Low	10	1	9
43	Error Code	Low	10	1	9
44	Crawler/Bot	Low	10	1	9
45	Anonymous Subscriber	Medium	50	25	25
46	Unidirectional Traffic	Low	10	5	5

Risk Info For Non-Experts

- nDPI not only detects issues but also tries to interpret data and explain to humans what is wrong with the analysed traffic.

Issues	Description
	Possible Exploit [Score: 250]
	TLS Cert. Mismatch [Score: 100] [?705810 vs hs-xen-rm-3,hs-xen-rm-3]
	TLS Cert. Self-signed [Score: 100] [CN=hs-xen-rm-3]
	Text With Non-Printable Chars [Score: 100] [?705810]
	Too Long TLS Cert. Validity [Score: 50] [TLS Cert lasts 3650 days]
	TLS not carrying HTTPS [Score: 10] [No ALPN]

Beyond Application Detection

- DPI toolkits were initially conceived as libraries able to detect the application protocol and extract metadata.
- Today most traffic is encrypted so we implemented ETA (Encrypted Traffic Analysis) and behaviour traffic analysis.
- nDPI over time also included algorithms and datatypes to provide network applications all it's needed to analyse traffic efficiently.

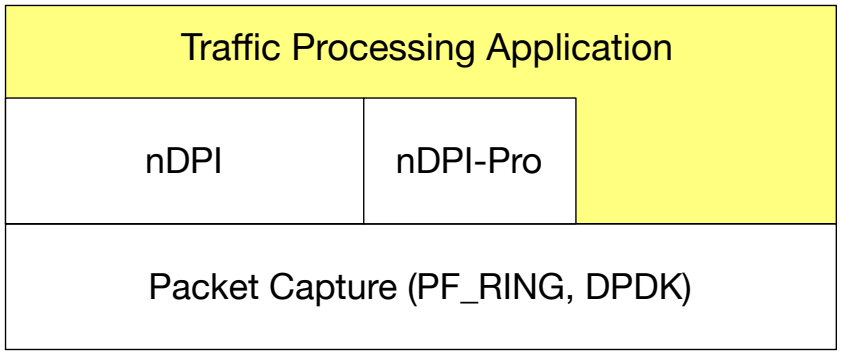

Detecting Malware

- Clear-text
  - Signatures <- too many signatures, slow.
  - Behaviour <- nDPI (e.g. binary application transfer)
- Encrypted traffic
  - Fingerprint and time/length bins (recognise encrypted traffic patterns)
  - Entropy (speculate about the content nature)

Malware Analysis: Trickbot

- Same packet sequence, same packet len and time distribution (using bins to detect similarities), same entropy…

```
00:08:02:1C:47:AE|20:E5:2A:B6:93:F1|0|0|0|10.9.25.101|5.53.125.13|49469|447|::|::|4|6|91|971|10|2732|11|1589100502|1589100502|27|
27|0.002|0.021|0.024|0|0|0|0|0|0||||||3,1,1,1,1,1|8,0,0,0,0,0|2.406|0.000|
95,-1419,134,-59,293,-773,-37,37,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
00:08:02:1C:47:AE|20:E5:2A:B6:93:F1|0|0|0|10.9.25.101|185.90.61.116|49482|447|::|::|4|6|91|931|9|2692|10|1589100502|1589100502|
27|27|0.002|0.025|0.026|0|0|0|0|0|0||||||3,1,1,1,1,1|8,0,0,0,0,0|2.406|0.000|
95,-1419,134,-59,293,-773,-37,37,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
00:08:02:1C:47:AE|20:E5:2A:B6:93:F1|0|0|0|10.9.25.101|195.123.221.104|49498|447|::|::|4|6|91|979|9|2692|10|1589100502|1589100502|
27|27|0.003|0.021|0.016|0|0|0|0|0|0||||||3,1,1,1,1,1|8,0,0,0,0,0|2.406|0.000|
127,-1419,134,-59,309,-773,-37,37,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0
00:08:02:1C:47:AE|20:E5:2A:B6:93:F1|0|0|0|10.9.25.101|195.123.221.178|49515|447|::|::|4|6|91|947|9|2692|10|1589100502|1589100502|
27|27|0.002|0.024|0.026|0|0|0|0|0|0||||||3,1,1,1,1,1|8,0,0,0,0,0|2.406|0.000|
95,-1419,134,-59,309,-773,-37,37,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,
```

Timeseries Analysis: Anomalies

- nDPI implements timeseries analysis logic to enable applications to detect anomalies without having to implement complex analysis techniques.

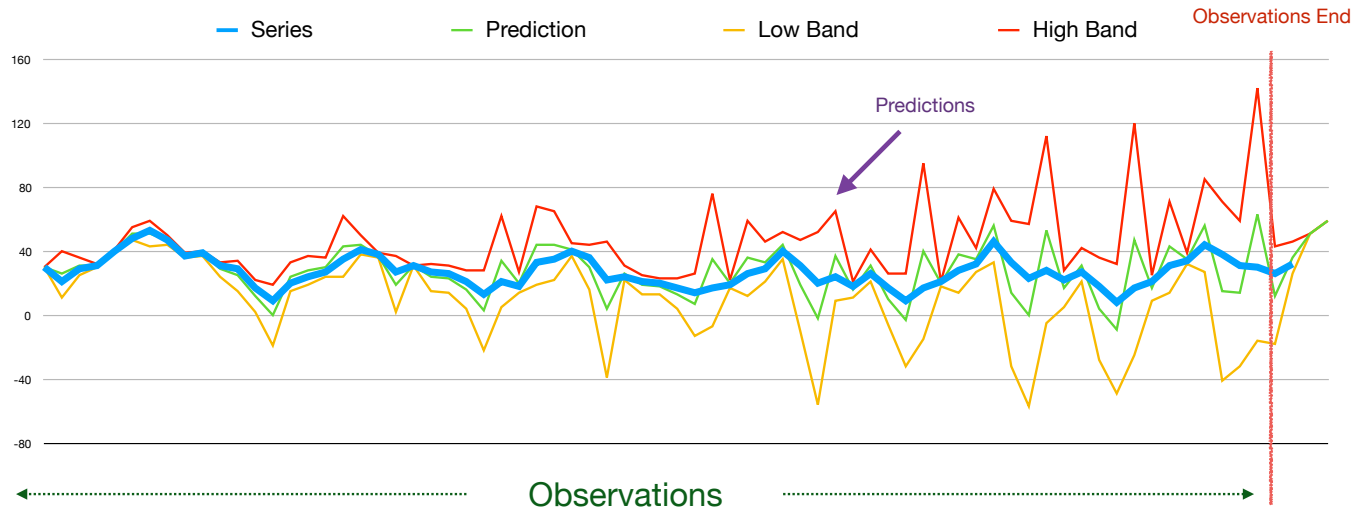

Timeseries Analysis: Similarity

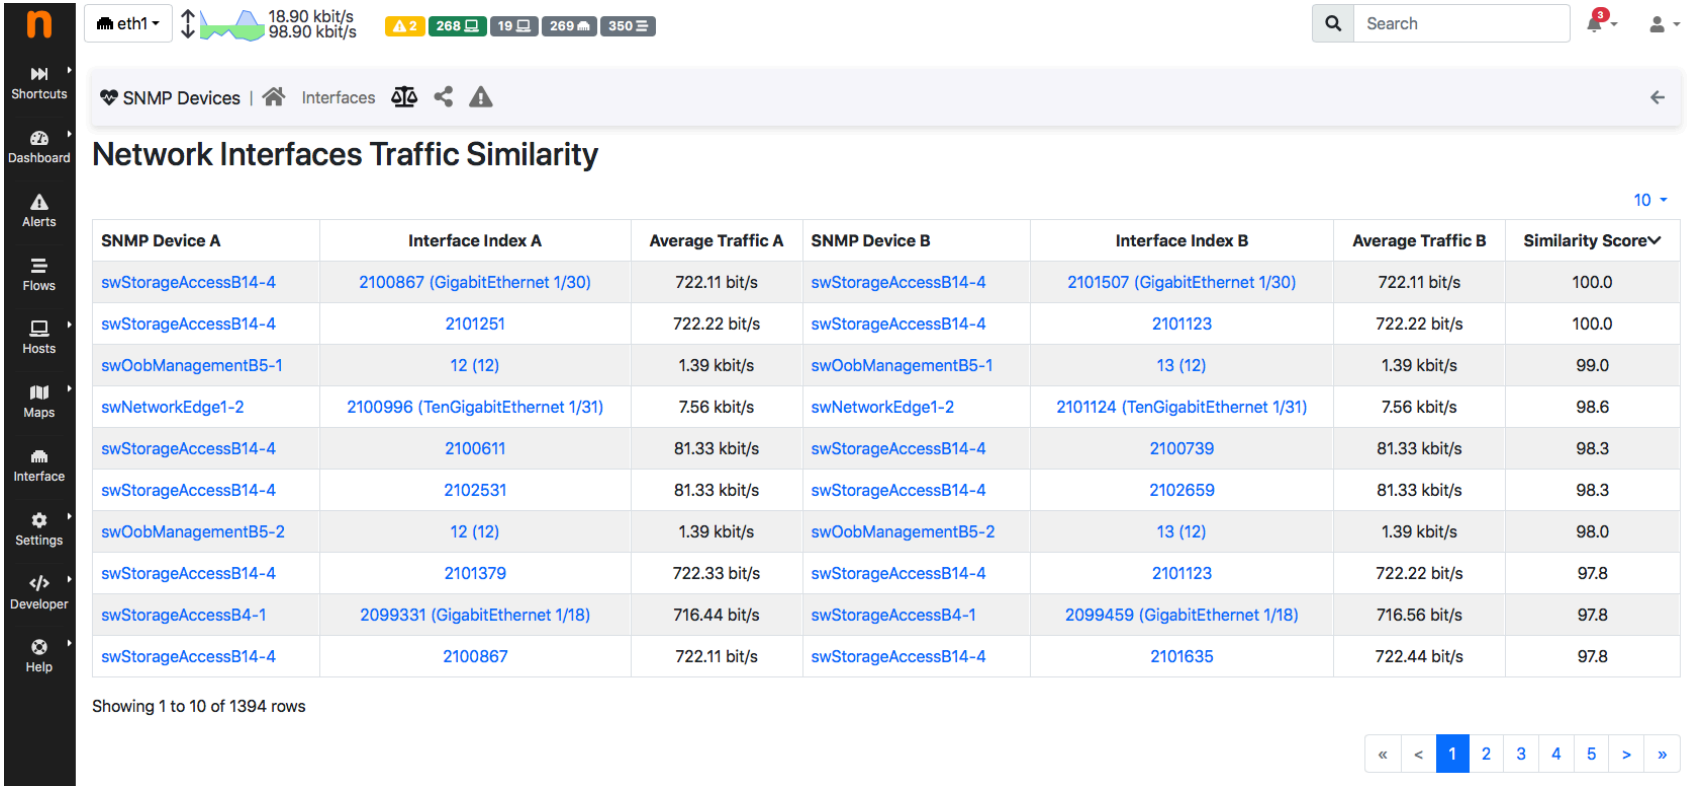

Network Interfaces Traffic Similarity

10

SNMP Device A	Interface Index A	Average Traffic A	SNMP Device B	Interface Index B	Average Traffic B	Similarity Score
swStorageAccessB14-4	2100867 (GigabitEthernet 1/30)	722.11 bit/s	swStorageAccessB14-4	2101507 (GigabitEthernet 1/30)	722.11 bit/s	100.0
swStorageAccessB14-4	2101251	722.22 bit/s	swStorageAccessB14-4	2101123	722.22 bit/s	100.0
swOobManagementB5-1	12 (12)	1.39 kbit/s	swOobManagementB5-1	13 (12)	1.39 kbit/s	99.0
swNetworkEdge1-2	2100996 (TenGigabitEthernet 1/31)	7.56 kbit/s	swNetworkEdge1-2	2101124 (TenGigabitEthernet 1/31)	7.56 kbit/s	98.6
swStorageAccessB14-4	2100611	81.33 kbit/s	swStorageAccessB14-4	2100739	81.33 kbit/s	98.3
swStorageAccessB14-4	2102531	81.33 kbit/s	swStorageAccessB14-4	2102659	81.33 kbit/s	98.3
swOobManagementB5-2	12 (12)	1.39 kbit/s	swOobManagementB5-2	13 (12)	1.39 kbit/s	98.0
swStorageAccessB14-4	2101379	722.33 bit/s	swStorageAccessB14-4	2101123	722.22 bit/s	97.8
swStorageAccessB4-1	2099331 (GigabitEthernet 1/18)	716.44 bit/s	swStorageAccessB4-1	2099459 (GigabitEthernet 1/18)	716.56 bit/s	97.8
swStorageAccessB14-4	2100867	722.11 bit/s	swStorageAccessB14-4	2101635	722.44 bit/s	97.8

Showing 1 to 10 of 1394 rows

« < 1 2 3 4 5 > »

Part II: nProbe

nProbe: A Recap

- nProbe is the oldest product ntop develops (since 2002).
- Initially conceived as a drop-in replacements for Cisco NetFlow analysers it is not a versatile tool for generating and emitting flow-based information.
- It is available in the “classic” nProbe (extremely versatile but not able to go above 10 Gbit) or “cento” designed for 40/100 Gbit networks (less versatile but speed native).
- Both products sit on top of nDPI.
- Support for Linux, Windows, MacOS, BSD (OPNsense, pfSense)

nProbe Cento

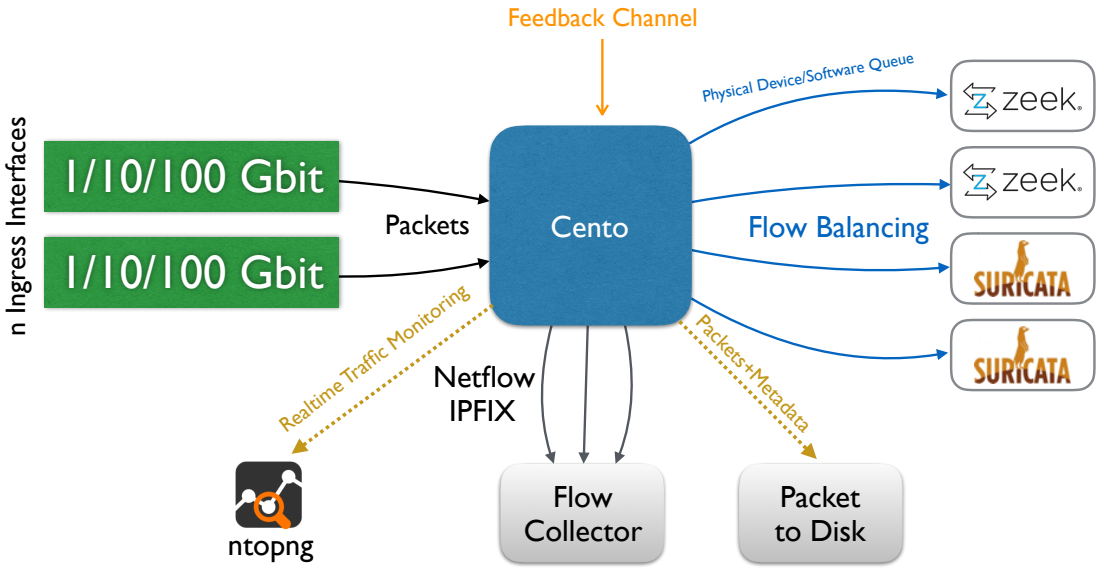

nProbe Recap: Flow Collection

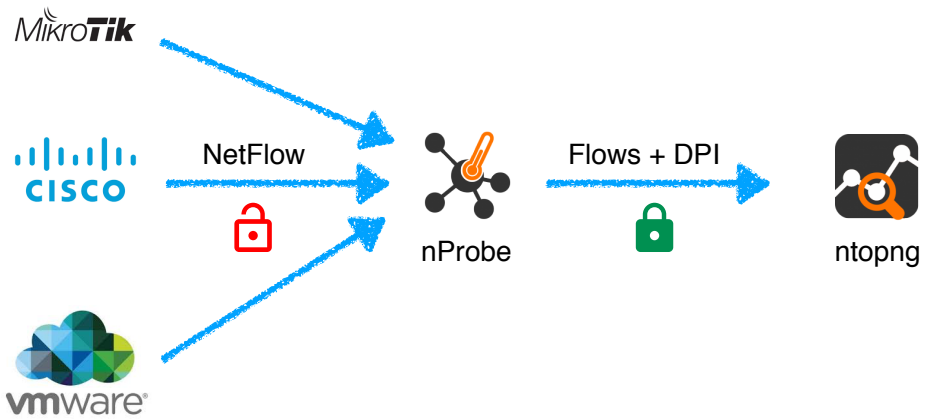

nprobe -3 <collector port> -i none -n none —zmq “tcp://*:1234” --zmq-encryption-key <pub key>

nProbe Recap: Flow Generation

nprobe -i tcp://localhost:1234 --zmq-encryption-key <pub key>

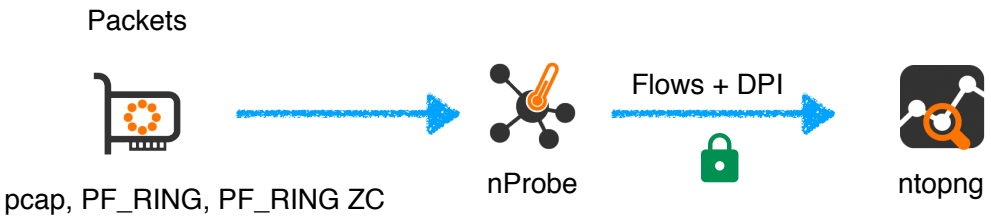

nprobe -i <interface> -n none --zmq tcp://localhost:1234 --zmq-encryption-key <pub key>

nProbe Recap: Datalakes

nprobe -i <device> …. --elastic <index type>;<index name>;<es URL>;<es user>:<es pwd>

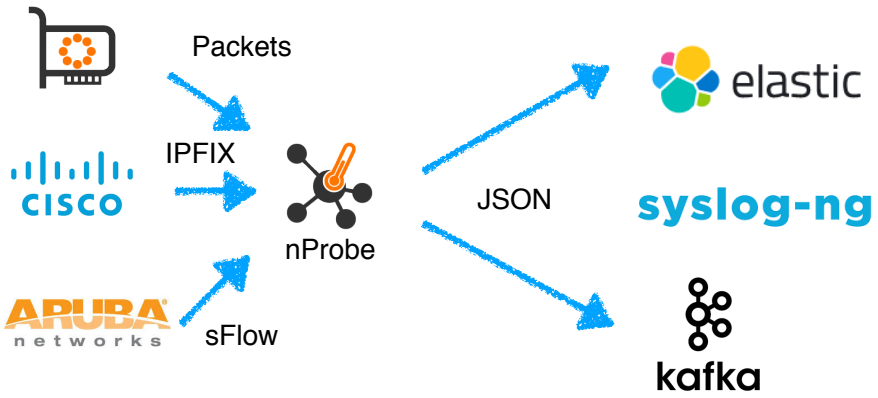

nprobe -i <device> …. --kafka <brokers>;<topic>;[<opt topic>;<ack>;<comp>]

Welcome to nProbe 10

- After a long development cycle, this July we will introduce a new major release that introduces several improvements over the 9.x series.
- Highlights:
  - Optimised flow collection for high-speed proprietary flow format handling.
  - Timeseries support: turn flows into InfluxDB timeseries.
  - Native monitoring of processes (Linux and Windows) and containers (Linux).
  - AWS VPC, Calix, Nokia AAA/NAT, native eBPF….

Proprietary NetFlow/IPFIX Collection [1/2]

- nProbe has been greatly enhanced to support proprietary information elements.
- Mapping happens via configuration files we share on github.
- nProbe maps fields at runtime by learning the type, length and name.

master ▾ nProbe / custom_fields /

lucaderi Added/redefined sgw-sgsnAddr and pgw-ggsnAddr

..	
AlcatelLucent	Reorganized custom fields
Calix	Added port
Cisco	Reorganized custom fields
Gigamon	Reorganized custom fields
Ixia	Reorganized custom fields
Nokia	Added/redefined sgw-sgsnAddr and pgw-ggsnAddr
PaloAlto	Renamed APPLICATION_NAME to APP_NAME to avoid name overlaps
Procera	Update procera_custom_fields.txt
Sonicwall	Updates sonicwall signature ids
Sophos	Adds Sophos AFC_FLOW_DIRECTION field
VMware	Added VMware custom fields
README.md	Updated URL

Proprietary NetFlow/IPFIX Collection [2/2]

```
# Name  STANDARD_ALIAS  PEN     FieldId         Len     Format

aaApp                       NONE    637     1           64      dump_as_ascii
aaAppGrp                    NONE    637     2           64      dump_as_ascii
aaSubType                   NONE    637     12          1       dump_as_uint
sessionDirection            NONE    637     13          1       dump_as_uint
natInsideSvcId              NONE    637     91          2       dump_as_uint
natOutsideSvcId             NONE    637     92          2       dump_as_uint
natSubString                NONE    637     93          64      dump_as_ascii
sessionDurationMilliseconds NONE    637     94              4       dump_as_uint
sessionStartSeconds         NONE    637     95              5       dump_as_uint
hostName                    NONE    637     96              64      dump_as_ascii
deviceId                    NONE    637     97              2       dump_as_uint
deviceMfgId                 NONE    637     98              2       dump_as_uint
deviceOsId                  NONE    637     99              2       dump_as_uint
deviceOsVer1                NONE    637     101              1       dump_as_uint
deviceOsVer2                NONE    637     102              1       dump_as_uint
deviceOsVer3                NONE    637     103              1       dump_as_uint
apn                         NONE    637     108              33      dump_as_ascii
mnc                         NONE    637     110              2       dump_as_uint
imsi                        NONE    637     111              8       dump_as_uint
msisdn                      NONE    637     112              8       dump_as_uint
sgw-sgsnAddr                NONE    637     113              16      dump_as_ipv6_address
pgw-ggsnAddr                NONE    637     114              16      dump_as_ipv6_address
ratType                     NONE    637     116              2       dump_as_uint
cellId                      NONE    637     118              4       dump_as_uint
imei                        NONE    637     129              8       dump_as_uint
mcc                         NONE    637     131              2       dump_as_uint
rlq                         NONE    637     321              1       dump_as_uint
```

nProbe: AWS Cloud

- Most cloud providers offer log-based monitoring tools for supervising hosted network activities.
- Options such as virtual TAPs are often expensive (BTW we have developed an OpenSwitch-based virtual taps in case you are interested) but logs are in proprietary formats preventing unified traffic monitoring.
- nProbe can now be used to turn logs into standard IPFIX flows

```
account-id action az-id bytes dstaddr dstport end flow-direction instance-id interface-id log-status packets pkt-dst-aws-service pkt-dstaddr pkt-src-aws-service pkt-srcaddr protocol region srcaddr srcport start sublocation-id
sublocation-type subnet-id tcp-flags traffic-path type version vpc-id
421717577885 ACCEPT use1-az6 396 10.113.39.219 80 1640154903 ingress - eni-0afec37a7c4be140d OK 5 - 10.113.39.219 - 10.113.39.208 6 us-east-1 10.113.39.208 7652 1640154859 - - subnet-048dbd0af4e64ae1f 3 - IPv4 5 vpc-0f4cdb08d3b1bcdf6
421717577885 ACCEPT use1-az6 1895 10.113.39.208 7652 1640154903 egress - eni-0afec37a7c4be140d OK 5 - 10.113.39.208 - 10.113.39.219 6 us-east-1 10.113.39.219 80 1640154859 - - subnet-048dbd0af4e64ae1f 19 1 IPv4 5 vpc-0f4cdb08d3b1bcdf6
421717577885 ACCEPT use1-az6 158 10.113.39.219 53540 1640154903 ingress - eni-0afec37a7c4be140d OK 1 - 10.113.39.219 - 10.112.84.16 17 us-east-1 10.112.84.16 53 1640154859 - - subnet-048dbd0af4e64ae1f 0 - IPv4 5 vpc-0f4cdb08d3b1bcdf6
421717577885 ACCEPT use1-az6 74 10.112.84.16 53 1640154903 egress - eni-0afec37a7c4be140d OK 1 - 10.112.84.16 - 10.113.39.219 17 us-east-1 10.113.39.219 53540 1640154859 - - subnet-048dbd0af4e64ae1f 0 1 IPv4 5 vpc-0f4cdb08d3b1bcdf6
421717577885 ACCEPT use1-az6 396 10.113.39.219 80 1640154903 ingress - eni-0afec37a7c4be140d OK 5 - 10.113.39.219 - 10.113.39.208 6 us-east-1 10.113.39.208 7568 1640154859 - - subnet-048dbd0af4e64ae1f 3 - IPv4 5 vpc-0f4cdb08d3b1bcdf6
421717577885 ACCEPT use1-az6 1895 10.113.39.208 7568 1640154903 egress - eni-0afec37a7c4be140d OK 5 - 10.113.39.208 - 10.113.39.219 6 us-east-1 10.113.39.219 80 1640154859 - - subnet-048dbd0af4e64ae1f 19 1 IPv4 5 vpc-0f4cdb08d3b1bcdf6
```

nProbe: Google Cloud

- For Google cloud users, we now offer the ability to push flows to Google Pub/Sub for creating cloud-based datalakes.
- nProbe can collect (or generate from packet capture) flows, convert them to JSON and push them to Google Cloud in addition to the typical consumers (e.g. ntopng).

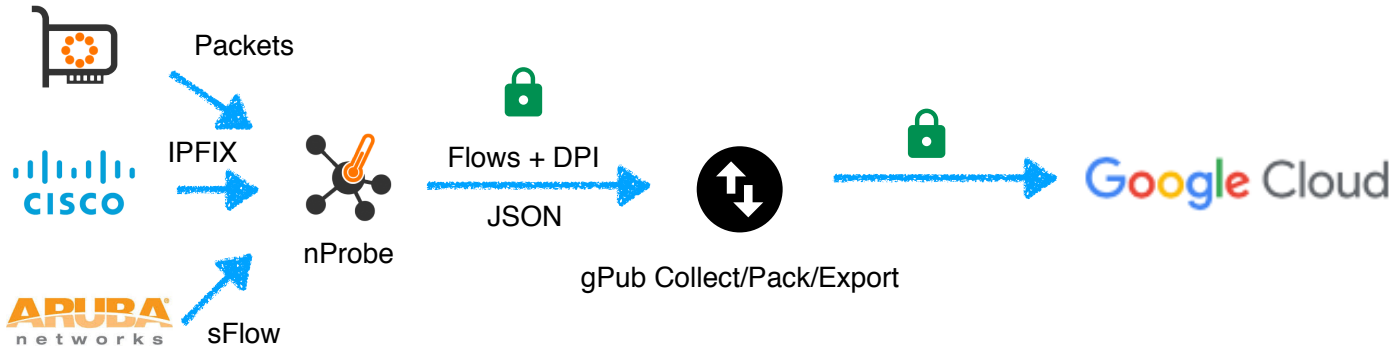

nProbe in Passive Optical Networks

- Modern broadband networks are based on passive optical networks (PON).
- nProbe has been enhanced to collect flows coming PONs are enable network operators to monitor their broadband networks.

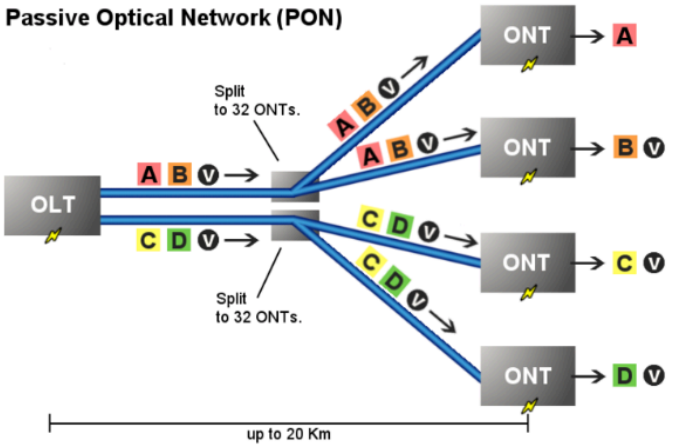

Source Wikipedia

```
> use calix;
Using database calix
> select * from calix_ipfix;
name: calix_ipfix
time                down_bytes  hostname ont_id    pon_ont_util_aid up_bytes
----                ----------  -------- ------    ---------------- --------
1638909102000000000 13980732227 OLT10    olt10test OLT10/1/1/gp8    44795119
1638909132000000000 13980733994 OLT10    olt10test OLT10/1/1/gp8    44798055
1638909162000000000 13980735371 OLT10    olt10test OLT10/1/1/gp8    44799898
1638909192000000000 13980737382 OLT10    olt10test OLT10/1/1/gp8    44803008
1638909222000000000 13980738425 OLT10    olt10test OLT10/1/1/gp8    44804390
1638909252000000000 13980742574 OLT10    olt10test OLT10/1/1/gp8    44810872
......
```

nProbe: Policy Enforcement [1/3]

- nProbe can be deployed in IPS mode on Linux (netfilter) and OPNsense/pfSense (netmap)

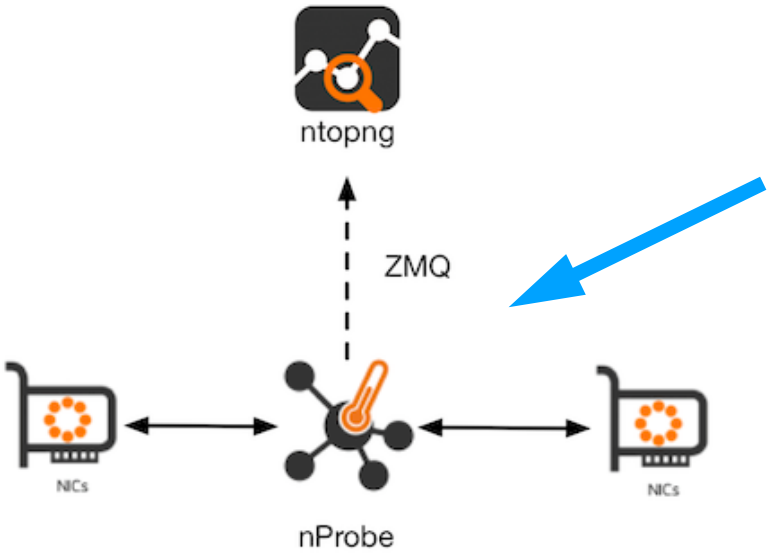

Typical deployment is close to the gateway (nord/sud traffic)

nProbe: Policy Enforcement [2/3]

- nDPI-pro implements and interpreter for a DPI-based enforcement engine.

```
### Pool definition ###

{ "pool": { "id":1, "name": "my pool 1", "ip":  [ "192.168.0.1/24", "10.0.0.0/8", "2a03:b0c0:2:d0::360:4001/48"], "mac": [] }, "policy": { "id":1 }}
{ "pool": { "id":2, "name": "my pool 2", "mac": ["e8:06:88:ff:fe:e4", "02:81:27:b5:f9:f3", "00:01:01:e4:ba:2c"],  "ip":  ["172.16.0.0/16"] }, "policy": { "id":2 }}
{ "pool": { "id":3, "name": "my pool 3", "ip":  [ "131.114.0.0/16" ], "mac": [] }, "policy": { "id": 3 }}

### Policy definition ###
# Continents: Africa / Asia-Pacific / Europe / North America / South America

# Root
{ "policy": { "id": 0, "name": "root rule", "default_marker": "pass", "flow_risk": { "risks": [ 12 ], "marker": "drop" }, "markers": { "categories": { "Video": "drop" }, "protocols":
{ "TLS": "pass" }}}}


# Rule definition (son of rule 0)
{ "policy": { "id": 1, "root": 0, "name": "my rule 1", "default_marker": "pass", "markers": { "protocols": { "HTTP": "pass" }, "countries": { "IT": "pass", "CN": "drop" },
"continents" : { "Asia" : "drop"}}}}
{ "policy": { "id": 3, "root": 0, "name": "my rule 3", "default_marker": "pass", "markers": { "protocols": { "HTTP": "drop" }}}}

# Subrule of rule 1 with more restrictions
{ "policy": { "id": 2, "root":1, "name": "my subrule 2 (son of rule 1)", "default_marker": "drop", "markers": { "protocols": { "53": "pass" }}, "hostnames":
{ "fundingchoicesmessages.google.com": "pass", "www.gstatic.com": "drop", "www.youtube.com": "pass" }}}


### Category files ###
{ "category_file": "../ndpi-pro/lists/nfw_malware_list.txt" }
{ "category_file": "../ndpi-pro/lists/nfw_mining_list.txt" }

### GeoIP ###
{ "geoip": { "asn": "../ndpi-pro/geoip/GeoLite2-ASN.mmdb", "city": "../ndpi-pro/geoip/GeoLite2-City.mmdb" }}
```

Complements Native Application Firewall Policing

nProbe: Policy Enforcement [3/3]

- With nProbe IPS, each host pool can have custom traffic policies configured by ntopng and enforced by nProbe IPS.

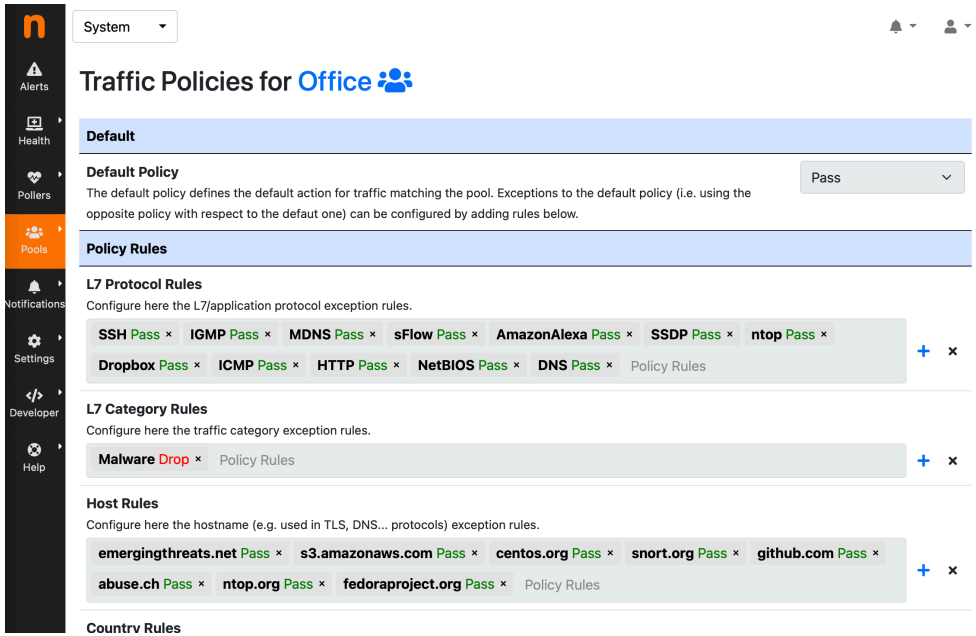

Introducing System Visibility in nProbe [1/5]

- nProbe:
  - Sits on top of the network stack (including containers) in order to receive traffic and inspect/block it.
  - Listen to system events in order to bind local traffic to processes and users.

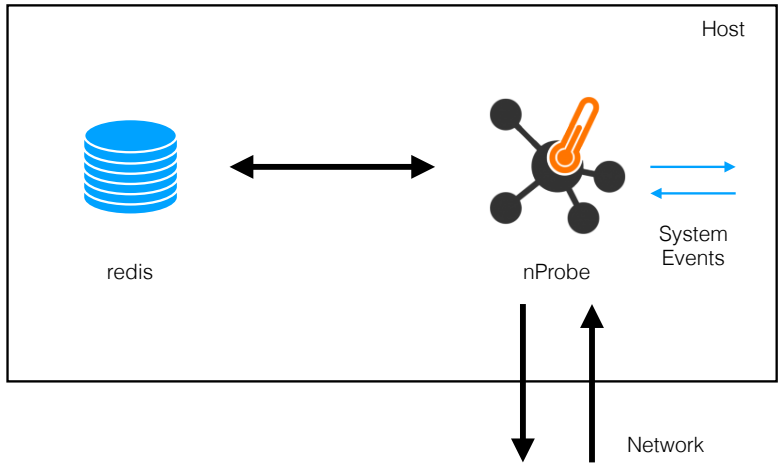

Introducing System Visibility in nProbe [2/5]

- nProbe uses redis as local policy cache for storing learnt information and as inter-process communication in case of high traffic rates that need to be handled by multiple nProbe processes.
- During the learning period, nProbe stores on redis observed <user>:<process> associations.
- Past learning, redis is used to retrieve known policies to be used for enforcement.

Introducing System Visibility in nProbe [3/5]

- Unless you are developing software, applications need to be installed with packages.
- Malware applications are (usually) not packaged, so this can be a good indicator of compromise.
- Currently we support Linux packaging: both .deb and .rpm families are supported.
- Windows is supported for app handling but not for packaging (not available).
- Alerts will be emitted for reporting the above issues.

Introducing System Visibility in nProbe [4/5]

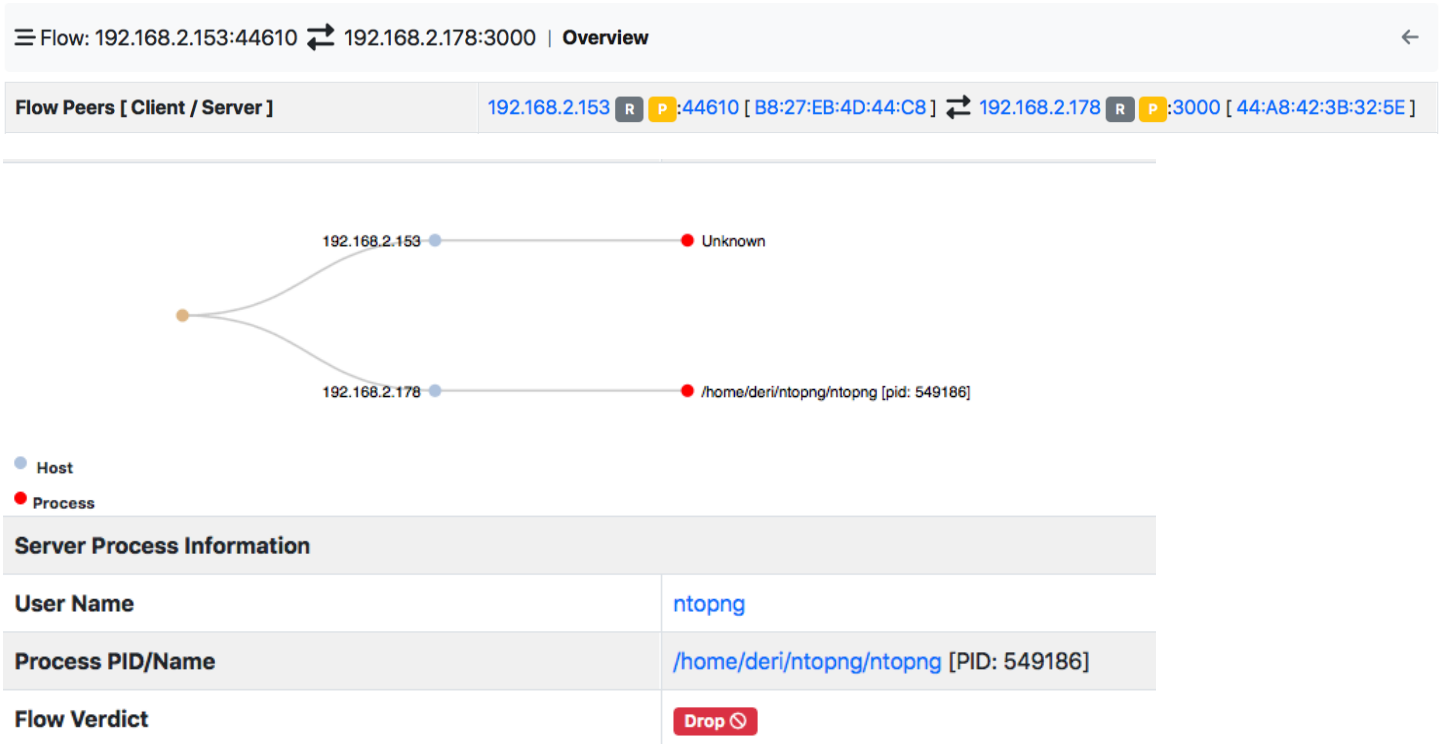


```
# nprobe -i eno1 --zmq tcp://127.0.0.1:1234 --agent-mode
```

```
# ntopng -i tcp://127.0.0.1:1234
```

Introducing System Visibility in nProbe [5/5]

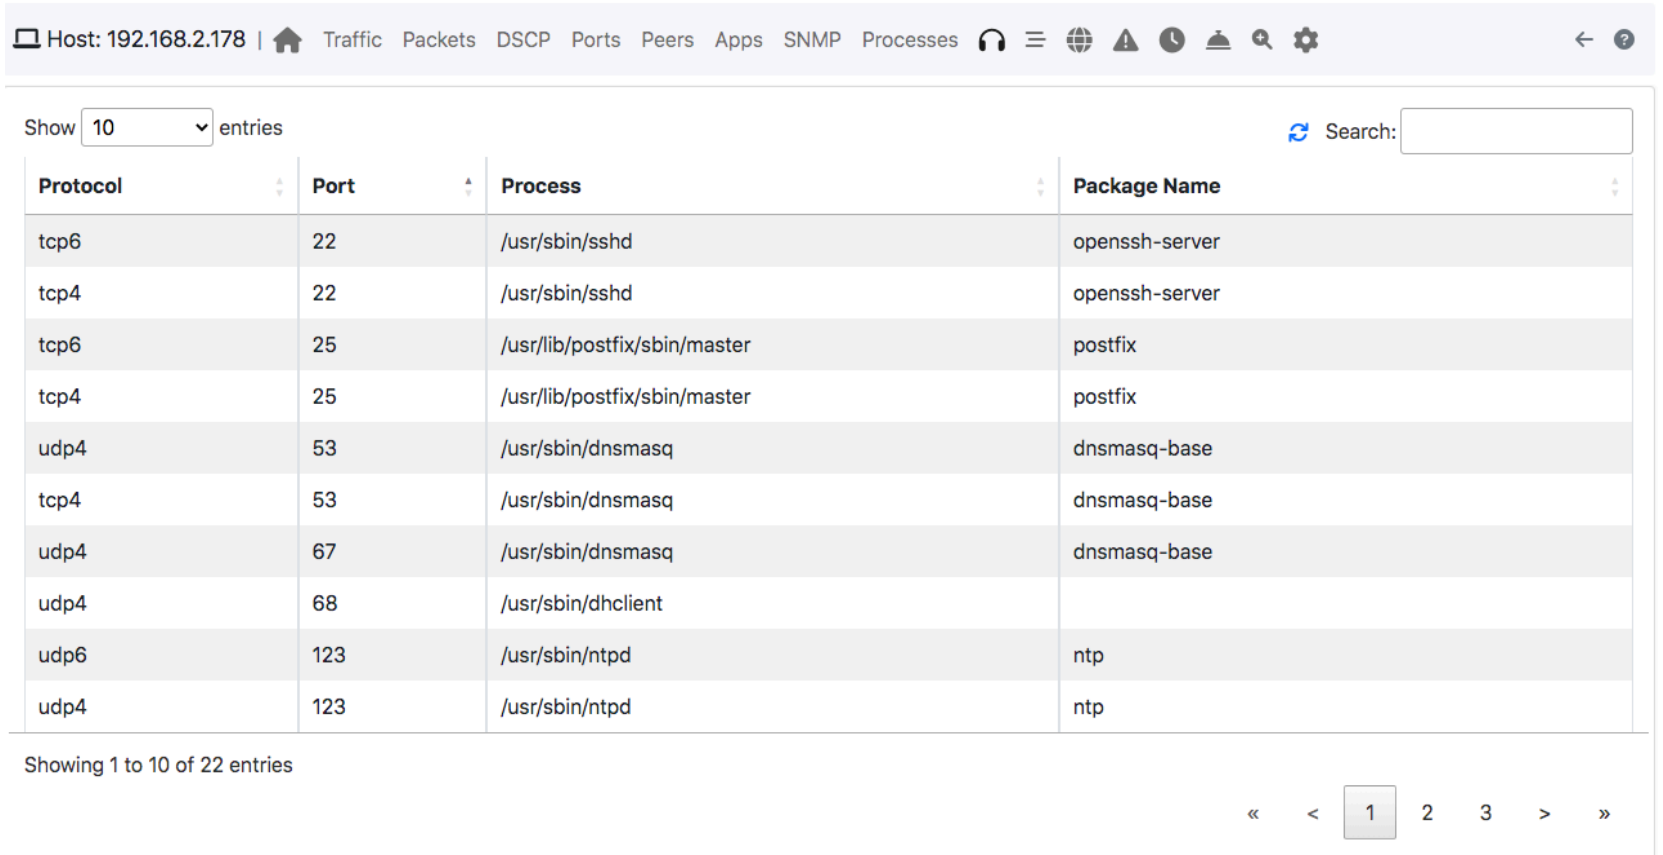

Summary

- nDPI evolved from a library for application protocol detection, to a comprehensive toolkit for traffic analysis supporting clear-text and encrypted protocols.
- nDPI-Pro created a layer over nDPI for instrumenting policies based on DPI metadata, as well provide process/container visibility based on eBPF (Linux) or native APIs (Windows).
- nProbe 10 introduces several new features including lightweight EDR facilities (whose metadata is exported in IPFIX/JSON), timeseries support, and custom information elements.

Joining ntop

jobs@ntop.org