

Cybersecurity at ntop: Present and Future

Luca Deri <deri@ntop.org>
@lucaderi



Part I: Past and Present

20+ Years of Network Monitoring

- Increased speed:
 - 10 Gbit is now commodity for many companies.
 - 100 Gbit is standard for ISPs.
- Monitoring Protocols
 - Still NetFlow and sFlow, just at higher speed.
 - Packet/Flow sampling prevents full visibility.
- Monitoring Metrics
 - Bytes and packets are still the main metrics for many network vendors.

IDSs and ML [1/2]

- Traditional IDS, often based on signatures and rule-based approaches shown their limitations in detection capability, especially when attackers heavily rely on encryption to obfuscate communications.
- While we do believe that ML (machine learning) technologies are playing (and will play in the future) an important role in cybersecurity, we strongly believe that domain knowledge and feature engineering have tremendous value for any detection problem.

IDSs and ML [2/2]

- Increasing adoption of encryption technologies, DPI (deep packet inspection) can be used to extract very strong signals from the raw traffic.
- While one could feed those signals to ML-based detectors, we highlight that when strong signals are available, one can greatly profit from them even with less sophisticated data processing technologies.
- This presentation shows how real-time, DPI-based cyber threat detection is feasible and effective using the concepts that will be explained later.

Signature-based IDSs (1998-Today)

```
alert tcp any any -> any [443,465] (msg:"Detected non-TLS  
on TLS port"; flow:to_server; app-layer-protocol:!tls;  
threshold: type limit, track by_src, seconds 90, count 1; sid:  
210003; rev:1;)
```

```
alert tcp any any <> any 443  
(msg:"APT.Backdoor.MSIL.SUNBURST"; content:"|16 03|";  
depth:2; content:"|55 04 03|"; distance:0;  
content:"digitalcollege.org"; within:50; sid:77600846; rev:1;)
```

- Techniques easy to circumvent.
- No application protocol visibility (packet header only, byte-based payload analysis).
- Outdated and error-prone format ("proto=TLS and SNI=digitalcollege.org").

Cybersecurity and Network Edge [1/2]

- Today most traffic is encrypted (80%+) and traditional clear-text protocols are moving to encryption (e.g. DNS vs DNS-over-HTTPS).
- As edge network speed is increasing, security threats on customer networks can propagate the issue to the core.
- Insecure devices (e.g. simple IoT devices) are placed in privileged network segments, thus requiring accurate supervision as they can cause severe troubles in case of breach.

Cybersecurity and Network Edge [2/2]

- Data centers with unhealthy customer traffic can affect neighbours and decrease the whole network reputation score.
- Limiting traffic observability to bandwidth usage is no longer wise: it is time to monitor customer traffic in an unobtrusive way in order to report users all threats they have not detected, mitigate issues and thus implement a healthier Internet.
- In essence we need to implement a lightweight (Raspberry an up, no GPU or GB of RAM) and scalable system able to model and analyse network traffic on a per-device basis, and being able to track device changes in behaviour.

Welcome to nDPI

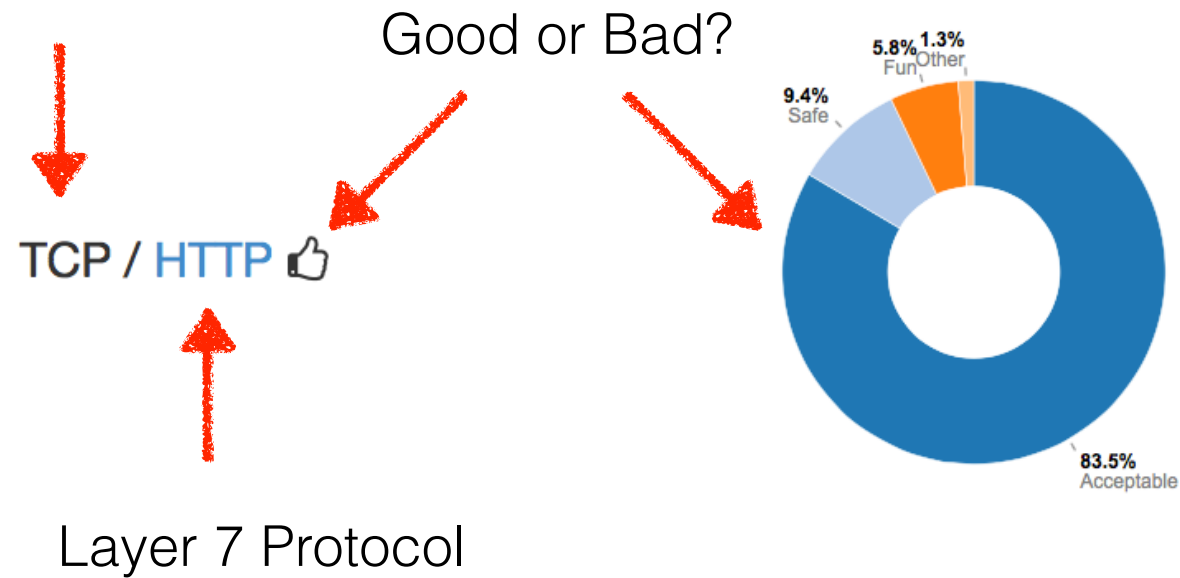
- In 2012 we decided to develop our own GNU LGPL DPI toolkit order to build an open source DPI layer.
- Protocols supported exceed 250+ and include:
 - P2P (BitTorrent)
 - Messaging (Viber, Whatsapp, Telegram, Facebook)
 - Multimedia (YouTube, Last.fm, iTunes)
 - Conferencing (Skype, Webex, Teams, Meet, Zoom)
 - Streaming (Zattoo, Disney, Netflix)
 - Business (VNC, RDP, Citrix)
 - Gaming



 **GitHub** <https://github.com/ntop/nDPI>

nDPI Traffic Analysis

Layer 4 Protocol



nDPI in Cybersecurity

- Analyses encrypted traffic to detect issues un-inspectable due to encrypted payload content.
- Extracts metadata from selected protocols (e.g. DNS, HTTP, TLS..) and matches it against known algorithms for detecting selected threats (e.g. DGA hosts, Domain Generated Algorithm).
- Associates a “flow risk” with specific flows to identify communications that are affected by security issues.

nDPI: Flow Risks

- HTTP suspicious user-agent
- HTTP numeric IP host contacted
- HTTP suspicious URL
- HTTP suspicious protocol header
- TLS connections not carrying HTTPS (e.g. a VPN over TLS)
- Suspicious DGA domain contacted
- Malformed packet
- SSH/SMB obsolete protocol or application version
- TLS suspicious ESNi usage
- Unsafe Protocol used
- Suspicious DNS traffic
- TLS with no SNI
- XSS (Cross Site Scripting)
- SQL Injection
- Arbitrary Code Injection/Execution
- Binary/.exe application transfer (e.g. in HTTP)
- Known protocol on non standard port
- TLS self-signed certificate
- TLS obsolete version
- TLS weak cipher
- TLS certificate expired
- TLS certificate mismatch
- DNS suspicious traffic
- HTTP suspicious content
- Risky ASN
- Risky Domain Name
- Malicious JA3 Fingerprint
- Malicious SHA1 Certificate
- Desktop of File Sharing Session
- TLS Uncommon ALPN
- TLS Certificate Validity Too Long
- Suspicious TLS Extension
- TLS Fatal Alert
- Suspicious Protocol traffic Entropy
- Clear-text Credentials Exchanged
- DNS Large Packet
- DNS Fragmented Traffic
- Invalid Characters Detected



Implemented Recently
(nDPI is a live project !)

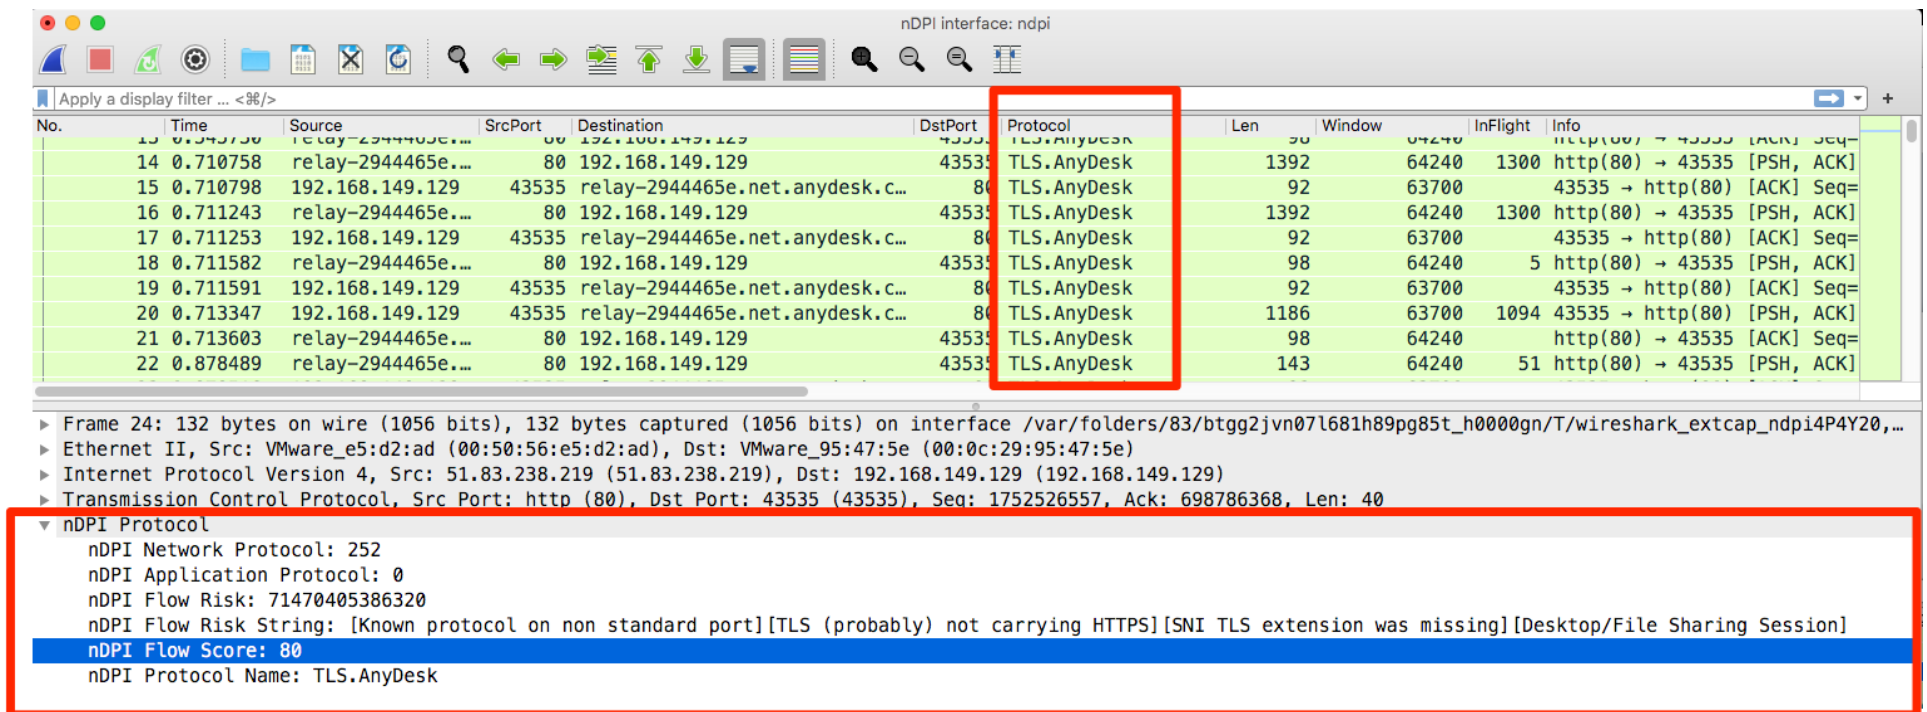
nDPI Encrypted Traffic Analysis

```
TCP 10.9.25.101:49184 <=> 187.58.56.26:449 [byte_dist_mean: 124.148883][byte_dist_std: 58.169660][entropy: 5.892724][total_entropy: 7124.302784][score: 0.9973][proto: 91/TLS][cat: Web/5][97 pkts/36053 bytes <=> 159 pkts/149429 bytes][Goodput ratio: 85/94][111.31 sec][bytes ratio: -0.611 (Download)][IAT c2s/s2c min/avg/max/stddev: 0/0 1129/662 19127/19233 2990/2294][Pkt Len c2s/s2c min/avg/max/stddev: 54/54 372/940 1514/1514 530/631][Risk: ** Self-signed Certificate **** Obsolete TLS version (< 1.1) **][TLsv1][JA3S: 623de93db17d313345d7ea481e7443cf][Issuer: C=AU, ST=Some-State, O=Internet Widgits Pty Ltd][Subject: C=AU, ST=Some-State, O=Internet Widgits Pty Ltd][Certificate SHA-1: DD:EB:4A:36:6A:2B:50:DA:5F:B5:DB:07:55:9A:92:B0:A3:52:5C:AD][Validity: 2019-07-23 10:32:39 - 2020-07-22 10:32:39][Cipher: TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA]
```

```
TCP 10.9.25.101:49165 <=> 144.91.69.195:80 [byte_dist_mean: 95.694525][byte_dist_std: 25.418150][entropy: 0.000000][total_entropy: 0.000000][score: 0.9943][proto: 7/HTTP][cat: Web/5][203 pkts/11127 bytes <=> 500 pkts/706336 bytes][Goodput ratio: 1/96][5.18 sec][Host: 144.91.69.195][bytes ratio: -0.969 (Download)][IAT c2s/s2c min/avg/max/stddev: 0/0 23/9 319/365 49/37][Pkt Len c2s/s2c min/avg/max/stddev: 54/54 55/1413 207/1514 11/134][URL: 144.91.69.195/solar.php][StatusCode: 200][ContentType: application/octet-stream][UserAgent: pwttyEKzNtGatwnJjmCcBLb0veCVpc][Risk: ** Binary application transfer **][PLAIN TEXT (GET /solar.php HTTP/1.1)]
```

Trickbot Traffic

nDPI in Wireshark



The screenshot displays the Wireshark network protocol analyzer interface. The top toolbar includes various icons for file operations, network analysis, and search. Below the toolbar, the 'Apply a display filter ... <%/>' field is visible. The main packet list pane shows a series of packets, with the 'Protocol' column highlighting 'TLS.AnyDesk' for several entries. A red box highlights the 'nDPI Protocol' details pane, which contains the following information:

- nDPI Network Protocol: 252
- nDPI Application Protocol: 0
- nDPI Flow Risk: 71470405386320
- nDPI Flow Risk String: [Known protocol on non standard port][TLS (probably) not carrying HTTPS][SNI TLS extension was missing][Desktop/File Sharing Session]
- nDPI Flow Score: 80**
- nDPI Protocol Name: TLS.AnyDesk

From Flow Risk To Score [1/2]

nDPI supported risks:

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 application transfer	Severe	250	125	125
5	Known protocol on non standard port	Medium	50	25	25
6	Self-signed Certificate	High	100	90	10
7	Obsolete TLS version (older than 1.2)	High	100	90	10
8	Weak TLS cipher	High	100	90	10
9	TLS Expired Certificate	High	100	50	50
10	TLS Certificate 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
15	TLS (probably) not carrying HTTPS	Low	10	5	5
16	Suspicious DGA domain name	High	100	90	10
17	Malformed packet	Low	10	5	5
18	SSH Obsolete Client Version/Cipher	High	100	90	10
19	SSH Obsolete Server Version/Cipher	Medium	50	5	45
20	SMB Insecure Version	High	100	90	10
21	TLS Suspicious ESNI Usage	Medium	50	25	25
22	Unsafe Protocol	Low	10	5	5
23	Suspicious DNS traffic	High	100	90	10
24	SNI TLS extension was missing	Medium	50	25	25
25	HTTP suspicious content	High	100	90	10
26	Risky ASN	Medium	50	25	25
27	Risky domain name	Medium	50	25	25
28	Possibly Malicious JA3 Fingerprint	Medium	50	25	25
29	Possibly Malicious SSL Cert. SHA1 Fingerprint	Medium	50	25	25
30	Desktop/File Sharing Session	Low	10	5	5
31	Uncommon TLS ALPN	Medium	50	25	25
32	TLS certificate validity longer than 13 months	Medium	50	25	25
33	TLS suspicious extension	High	100	90	10
34	TLS fatal alert	Low	10	5	5
35	Suspicious entropy	Medium	50	25	25
36	Clear-text credentials	High	100	90	10
37	DNS packet larger than 512 bytes	Medium	50	25	25
38	Fragmented DNS message	Medium	50	25	25
39	Text contains non-printable characters	High	100	90	10

From Flow Risk To Score [2/2]


[illegible]

Detected Risk	Risk Score Value
Known protocol on non standard port	10
TLS (probably) not carrying HTTPS	10
SNI TLS extension was missing	50
Desktop/File Sharing Session	10
Flow Score Total	80

Consolidating Score [1/3]

- Flow traffic analysis is too granular and it needs to be consolidated into:
 - Network Interface
 - Host/Network/Customer.
 - ASN/Country
- In essence that is the pillar for creating a (client/server) numerical score that can be quickly used to spot issues (network, security...).

Consolidating Score [3/3]



- Flow score is computed in realtime (flow lifetime)
- (Host/Interface/....) Checks are performed every minute


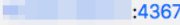

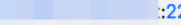





What about Risk Exceptions ? [1/3]

- Many cybersecurity products are very strict with policies and they divide the world in good and bad.
- Unfortunately reality is a bit more complicated (indeed grey exists), and “modern” needs to coexist with “ancient” that in computing terms can be just a few years old.
- The score principle is effective only if there are no false positives as otherwise they can deceive detection algorithms by generating false alerts.




What about Risk Exceptions ? [2/3]

- A few typical exception examples:
 - Private IPs with self-signed TLS certificates.
 - Insecure protocols/hosts that cannot be upgraded but that provide a specific service to a few clients.
 - Applications running on non standard ports (e.g. SSH server on port 2222).
 - TLS towards numeric IP address (no symbolic hostname).

What about Risk Exceptions ? [3/3]

Date/Time ↑↓	Score ↑↓	Application	Alert	Flow	Actions
 15:17:43	60	TCP:SSH	Obsolete SSH	 :43670   :22 	   



Available options:



- Disable Check (for everybody). 
- Exclude the check for a specific host. 
- Acknowledge the alert 

Exclude Checks: Obsolete SSH

Exclude Checks "Obsolete SSH". Exclude For:

☒ Any host (disable check)

☐  (,150)

☐  (,111)

Stored alerts matching the specified disable criteria be deleted.

☒ Delete Alerts

Checks matching the specified exclusion criteria will not be run and alerts will not be triggered.

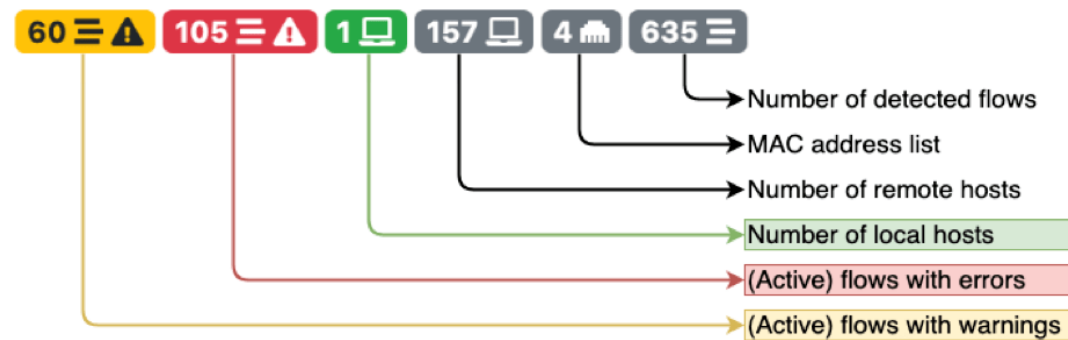
Exclude

Score At Work

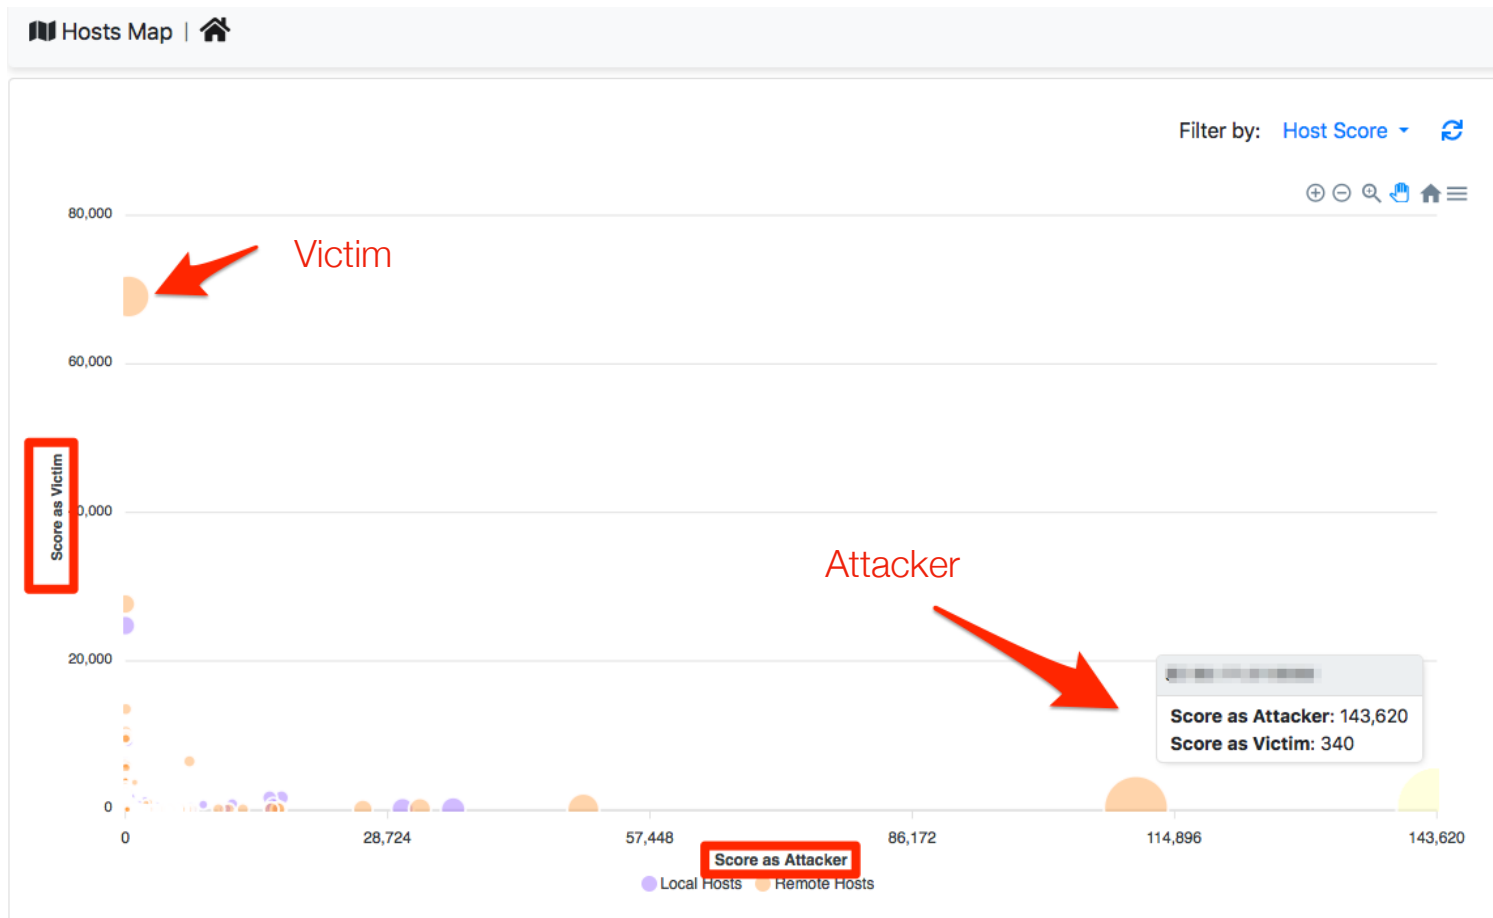


All Hosts

	IP Address	VLAN	Flows	Score	Name	Seen Since	Breakdown	Throughput	Total Bytes
	[IP] [US] R	250	9853	111,320	[Name]	03:19	Sent R	34.75 kbit/s ↑	642.7 KB
	[IP] [US] R	250	10854	102,850	[Name]	09:44:37	Sent R	47.07 kbit/s ↑	168.32 MB
	[IP] [RU] R	250	2231	73,815	[Name]	09:44:04	Rcvd	18.98 kbit/s ↑	64.26 MB
	[IP] [RU] R	250	823	52,938	[Name]	09:44:03	Sent Rcv	4.03 kbit/s ↓	21.5 MB

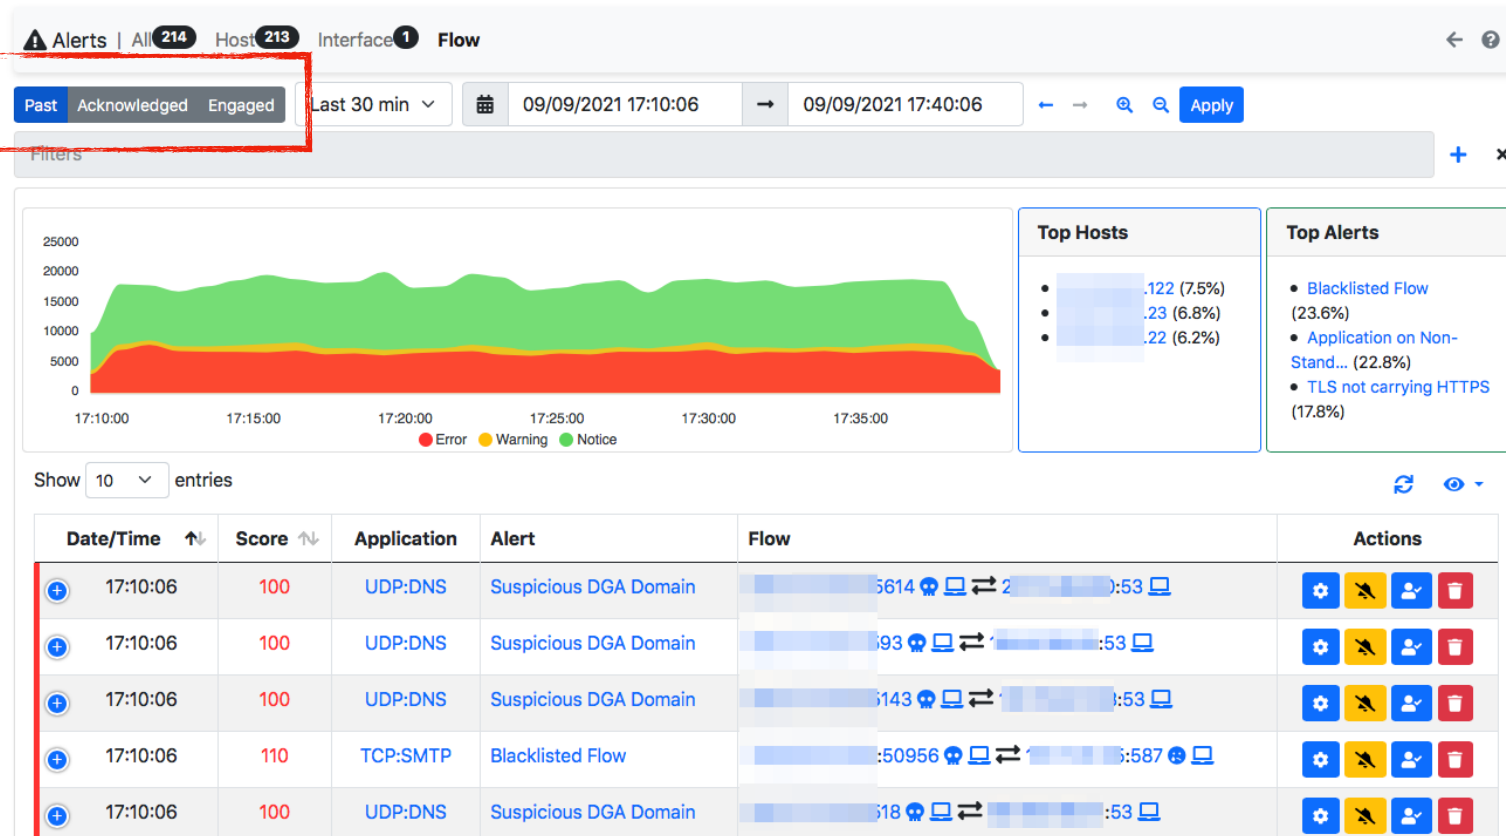


Visualising Cybersecurity: Bubbles



Score-based Alerts [1/2]

Alarm
Lifecycle



Score-based Alerts [2/2]

Flow score		Attacker		Victim	
Date/Time	Score	Application	Alert	Flow	Actions
17:10:06	100	UDP:DNS	Suspicious DGA Domain	:45614 → :53	
Description Suspicious DGA Domain [fdfb81ebb7184b4c6d3206117ad2531.ix.dnsbl.manitu.net]					
17:10:06	110	TCP:SMTP	Blacklisted Flow	:50956 → :587	
Description Blacklisted Client					
Other Issues Application on Non-Standard Port [Score: 10]					

Threshold-based Score Alerts [1/2]

Checks | Host | Interface | Local Network | SNMP Device | Flow | System | Syslog

All (16) Enabled (1) Disabled (15)

Filter Categories Search Script: score

Name	Category	Description	Values	Action
Dangerous Host		Trigger an alert when an host crosses the configured score threshold for more than 5 consecutive minutes	> 1000 Score (Minute)	
Score Threshold Exceeded		Trigger an alert when the score of an host exceeds the threshold	> 5000 Score (Minute)	

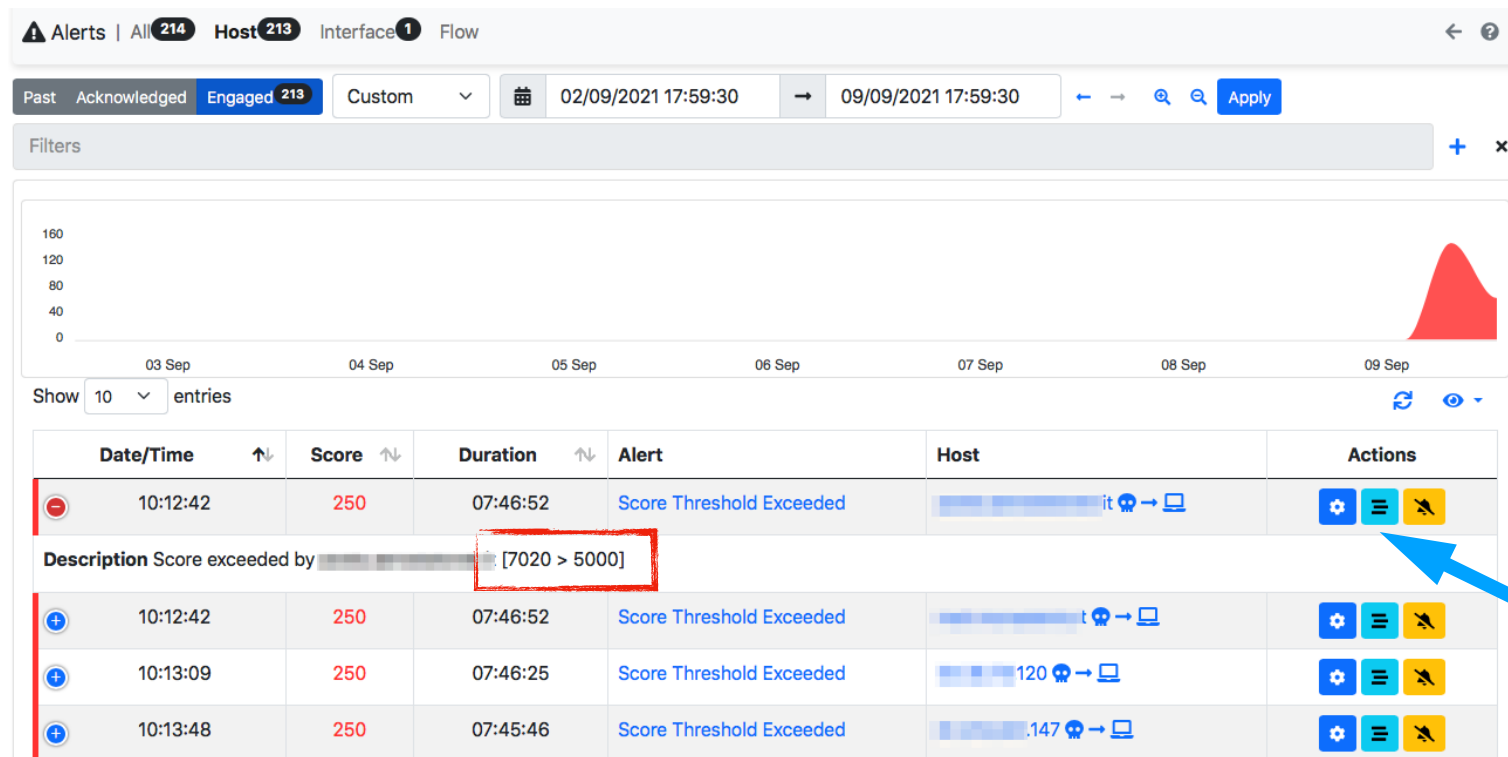
Showing 1 to 3 of 3 rows

« < 1 > »

Simple to use for detecting hosts with high score:

- Continuously
- Score spikes

Threshold-based Score Alerts [2/2]




Score-based Behaviour Analysis [1/5]

- Thresholds are useful to spot issues that can be identified with boundaries.
- However
 - How do you define a typical host threshold? Not all hosts behave the same way.
 - How can I detect changes in behaviour? A host can double its score and still be unalarmed, but the network operator needs to be informed that something has changed.

Score-based Behaviour Analysis [2/5]

- Without having to disturb ML that can be heavy for many users, we have decided to use (mature) statistical methods for spotting these changes.
- The advantage of statistical methods is that we can create a lightweight model per metric (hosts have tent of metrics) that uses little memory and CPU.
- For the record, we have used DES (Double Exponential Smoothing) that implements data forecasting and high/lower band for detecting changes in behaviour.

Score-based Behaviour Analysis [3/5]



Score-based Behaviour Analysis [4/5]

Checks | Host | Interface | Local Network | SNMP Device | Flow | System | Syslog

All (16) Enabled (1) Disabled (15)


Filter Categories Search Script: score

Name	Category	Description	Values	Action
Score Anomaly		Detects anomalies in host score		

Showing 1 to 3 of 3 rows

<< < 1 > >>

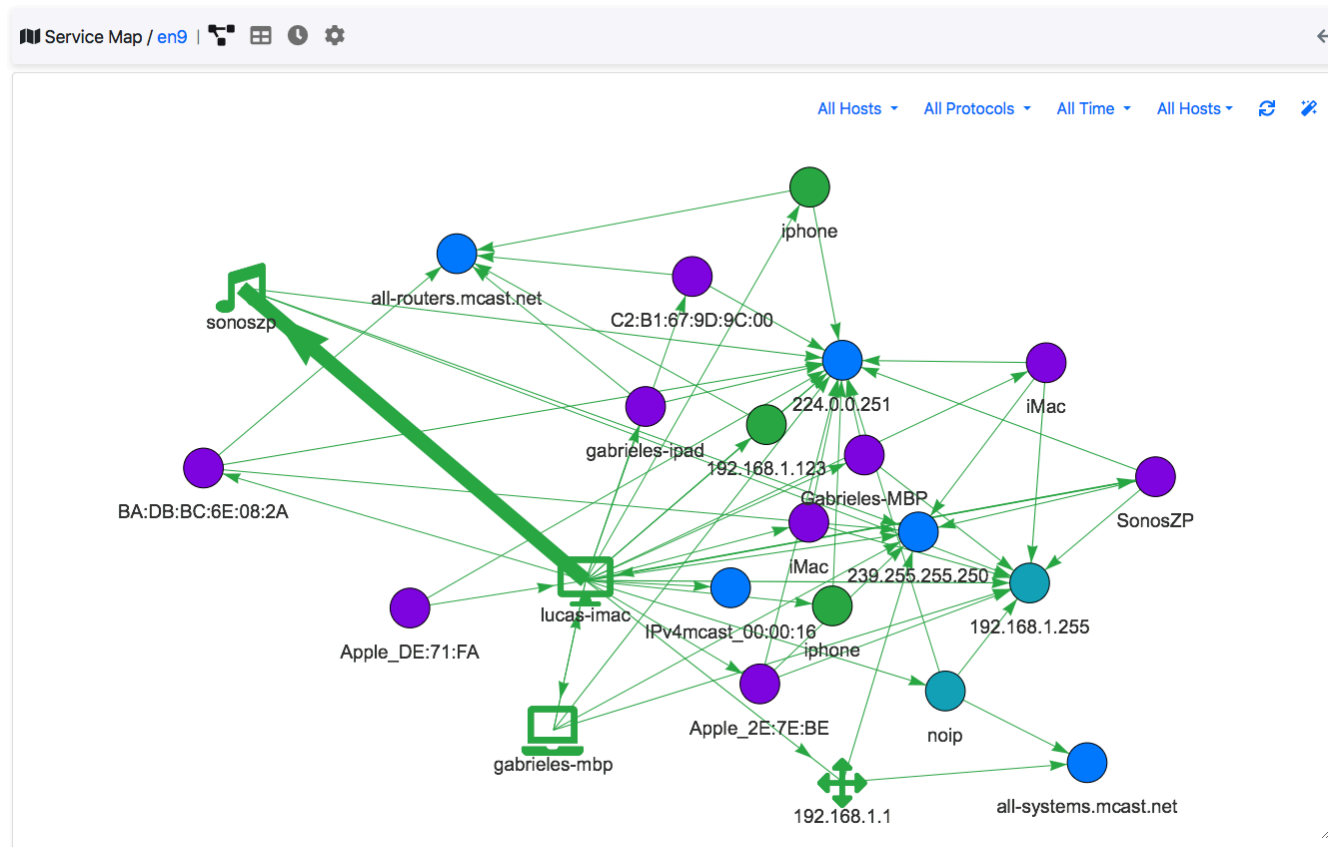
Score-based Behaviour Analysis [5/5]



Lateral Movement [1/4]

- What happens if a malware is roaming in our network?
How can we spot it?
- In addition to the checks just presented, it can help to create a model of the network traffic and to continuously match it against live communications.
- Communications not matching the model are probably an indication of mistakes or new traffic patterns worth to be analysed.


Lateral Movement [2/4]



Lateral Movement [3/4]

- Learning Period
 - Discover new services and assign a default policy to them.
 - No alert is generated during learning.

Post Learning
Alerts Enabled



Learning Period
Configure the learning period for behavioural traffic analysis.

Hours Days

Service Status During Learning
The default status of a new discovered service when the Service Map is learning.

Undecided Allowed Denied

Service Status Post Learning
The default status of a new discovered service when the Service Map has finished the learning.

Undecided Allowed Denied

Lateral Movement [4/4]

Service Map / en9 | [Icons]

Show 10 entries

All Protocols All Time Status Search: [Input]

Protocol	Client	Server	VLAN	Port	Contacts	Last Seen	Info	Service Status
UDP:MDNS	iMac	224.0.0.251	0	5353	77	01:38:27 ago	_spotify-connect_tcp.local	[Hourglass] [Checkmark] [X] Forbidden
UDP:MDNS	lucas-imac	iMac	0	5353	1	02:25:07 ago	luca__s_imac_companion-link_tcp.local	[Hourglass] [Checkmark] [X]
UDP:MDNS	C2:B1:67:9D:9C:00	224.0.0.251	0	5353	5	21 Days, 03:17:30 ago	_airplay_tcp.local	[Hourglass] [Checkmark] [X]
UDP:MDNS	50-35-10-70.1	224.0.0.251	0	5353	42	03:13 ago	1_airport_tcp.local	[Hourglass] [Checkmark] [X]
UDP:MDNS	iphone	224.0.0.251	0	5353	79	02:08 ago	_companion-link_tcp.local	[Hourglass] [Checkmark] [X]
UDP:MDNS	Apple_2E:7E:BE	224.0.0.251	0	5353	10	21 Days, 02:59:22 ago	macbook_companion-link_tcp.local	[Hourglass] [Checkmark] [X]
UDP:MDNS	iPhone	224.0.0.251	0	5353	64	09:53 ago	_sleep-proxy_udp.local	[Hourglass] [Checkmark] [X]
UDP:MDNS	lucas-imac	Apple_2E:7E:BE	0	5353	7	21 Days, 03:15:30 ago	_companion-link_tcp.local	[Hourglass] [Checkmark] [X]
UDP:MDNS	lucas-imac	Gabrieles-MBP	0	5353	1	03:23:50 ago	_companion-link_tcp.local	[Hourglass] [Checkmark] [X]
UDP:MDNS	lucas-imac	gabrieles-mbp	0	5353	19	03:14 ago	_smb_tcp.local	[Hourglass] [Checkmark] [X]

Showing 41 to 50 of 396 rows

« < 1 ... 4 5 6 ... 40 > »





Beaconing Detection [1/3]


- Beacons are periodic low-volume communications that can be easily hidden inside the overall traffic.
- They are:
 - Often used by malware to talk back with the master.
 - An indication of failures (e.g. periodic connection to a service that is unavailable).
 - Used to identify monitoring activities (e.g. scans etc) or periodic checks (e.g. email download).
- In essence beaconing is not just for cybersecurity but also for spotting activities worth to be analysed.

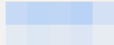
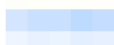
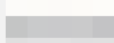

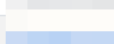
Beaconing Detection [2/3]

- Instead of using AI or complex algorithms for beaconing detection we use a simple method:
 - Keep track of quadruplets <source/destination IP, destination port, layer 4 protocol>.
 - As soon as a new flow is detected a quadruplet is created (if not already present) or updated (if already created).
 - Idle quadruplets or quadruplets whose periodicity isn't too constant (of course we take into account time drifts) are discarded.

Beaconing Detection [3/3]

Periodicity Map / 192.168.1.178 |    

Show 10 entries Protocol All Time Search: 

Protocol	Client	Server	Port	Observations	Frequency	Last Seen	Info
ICMP	Luca's iMac			144	3 sec	00:02 ago	
TCP:Google	Luca's iMac		4070	3	120 sec	00:33 ago	
TCP:IMAPS	Luca's iMac		993	3	120 sec	01:04 ago	
TCP:IMAPS	Luca's iMac		993	3	121 sec	01:03 ago	
TCP:IMAPS	Luca's iMac		993	3	120 sec	01:04 ago	

- Beaconing with Unknown or “unpleasant” (e.g. IRC) protocols are an indicator of suspicious communications.
- Beaconing begin/end is reported as informative alert.

Alerts: Actionable Items

The screenshot displays the ntopng web interface for managing alert endpoints. A modal window titled "Add New Endpoint" is open, showing a form with fields for "Name", "Type", and "WebHook URL". The "Type" dropdown menu is open, listing various alert types: Discord (selected), Elasticsearch, Email, Fail2Ban, Shell Script, Slack, Syslog, Telegram, and Webhook. Below the form, there are instructions for setting up a Discord webhook.

Add New Endpoint

Name

Type

- ✓ Discord
- Elasticsearch
- Email
- Fail2Ban
- Shell Script
- Slack
- Syslog
- Telegram
- Webhook

WebHook URL

Instructions:

- Open the Discord channel you want to receive alerts in.
- From the channel menu, click on Webhooks menu.
- Click on Webhooks menu.
- Click the Create Webhook button and note the name of the bot that will post the messages (note that you can set it on the ntopng recipients page)
- Note the URL from the WebHook URL field to be copied in the field above.
- Click the Save button.

Add

Background Interface:

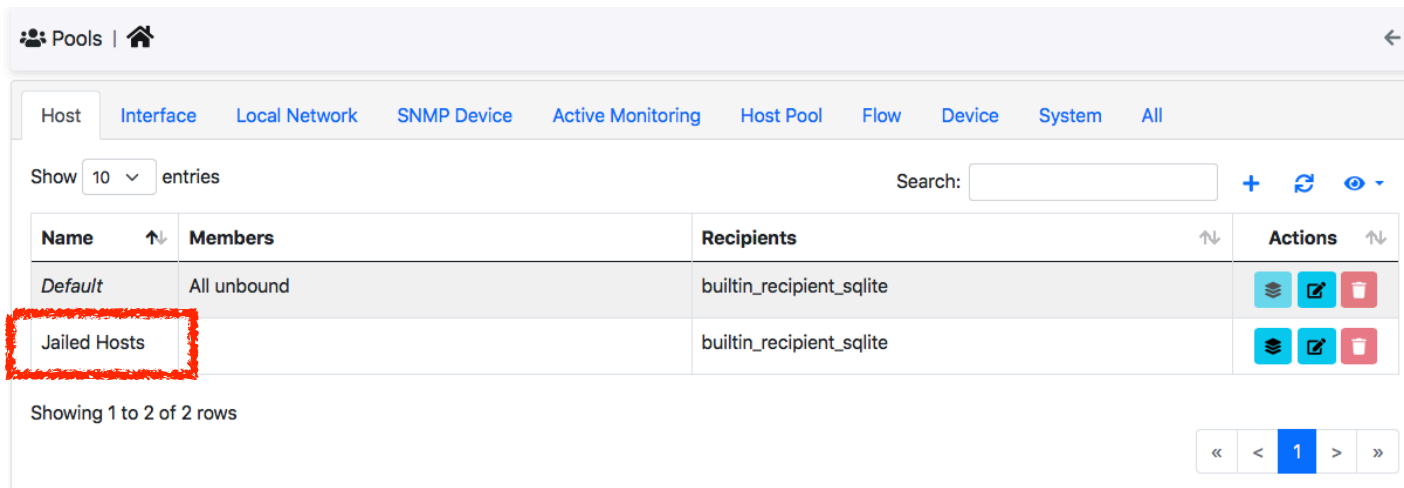
- System dropdown
- Endpoints | Home icon
- Show 10 entries
- Type Search:
- Table with columns: Name, Type, Used by Recipients, Actions
- Table rows: ban (Fail2Ban), builtin_endpoint_sqlite (SQLite built-in)
- Showing 1 to 2 of 2 rows
- Manage Configuration button
- ntopng Enterprise L v.5.1.210826 (M)
- Status bar: 09:28:30 +0200 | Uptime: 13:23

Using Score to Enforce Policies [1/7]







- A host pool is a logical group of hosts that for some reason (i.e. they do not need to belong to the same IP network or VLAN) can be grouped together.
- Pools can have alert actions defined: this allows hosts to perform different actions when an alert is triggered.
- Example:
 - Send a slack message to XYZ when there is alert for pool ABC.
 - Just log the alert for hosts other than XYZ.

Using Score to Enforce Policies [2/7]

- All pools are alike with the exception of the “Jailed Hosts” Pool.



The screenshot shows the ntop web interface for managing Host Pools. The breadcrumb navigation at the top reads "Pools | Home". Below this is a tabbed interface with tabs for Host, Interface, Local Network, SNMP Device, Active Monitoring, Host Pool, Flow, Device, System, and All. The "Host Pool" tab is selected. Under this tab, there is a "Show 10 entries" dropdown and a search bar. A table displays the list of pools. The "Jailed Hosts" pool is highlighted with a red dashed box. The table has columns for Name, Members, Recipients, and Actions. The "Default" pool has "All unbound" members and "builtin_recipient_sqlite" recipients. The "Jailed Hosts" pool also has "builtin_recipient_sqlite" recipients. The "Actions" column contains icons for adding, editing, and deleting a pool. At the bottom, it says "Showing 1 to 2 of 2 rows" and a pagination control shows page 1 of 1.

Name	Members	Recipients	Actions
Default	All unbound	builtin_recipient_sqlite	  
Jailed Hosts		builtin_recipient_sqlite	  

- Dangerous hosts are added/removed to/from this pool as they are detected or come back to normal.

Using Score to Enforce Policies [3/7]

- It is possible to use the host score in order to prevent hosts from “infecting” the rest of the network.

Trigger Family

Checks | All | Host | Interface | Local Network | SNMP Device | Flow | System | Syslog

All (121) Enabled (84) Disabled (37)

Intrusion Detection and Prevention Search Script:

Name	Category	Description	Values	Action
Dangerous Host		Triggers an alert and adds the host to the jailed hosts pool for 30 minutes, when the configured score threshold is cros...	> 1000 Score (Minute)	

Showing 1 to 1 of 1 rows

<< < 1 > >>

Stateful Alert (no permanent block)

Using Score to Enforce Policies [4/7]

Dangerous Host

Enabled

Minute

>

1000

Score

Excluded Hosts

Comma separated list of IP addresses. This alert won't be triggered for hosts inside this list.


Triggers an alert and adds the host to the jailed hosts pool for 30 minutes, when the configured score threshold is crossed for more than 5 consecutive minutes. Hosts in the jailed hosts pool are prevented from generating traffic when using nProbe in IPS mode.

Reset to Factory Value

Apply

Using Score to Enforce Policies [5/7]

- It is possible to combine ntop tools to enforce policies using nProbe in IPS mode:



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

Using Score to Enforce Policies [6/7]

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

The screenshot displays the ntopng web interface for configuring traffic policies. The left sidebar contains navigation links: Alerts, Health, Pollers, Pools (highlighted), Notifications, Settings, Developer, and Help. The main content area is titled 'Traffic Policies for Office' and shows a 'System' dropdown menu. Below the title, there's a 'Default' section with a 'Default Policy' dropdown set to 'Pass'. A description explains that the default policy defines the default action for traffic matching the pool, and exceptions can be configured by adding rules below. The 'Policy Rules' section is divided into three categories: L7 Protocol Rules, L7 Category Rules, and Host Rules. Each category has a description and a list of rules. L7 Protocol Rules include SSH, IGMP, MDNS, sFlow, AmazonAlexa, SSDP, ntop, Dropbox, ICMP, HTTP, NetBIOS, and DNS. L7 Category Rules include Malware. Host Rules include emergingthreats.net, s3.amazonaws.com, centos.org, snort.org, github.com, abuse.ch, ntop.org, and fedoraproject.org. Each rule has a 'Pass' or 'Drop' status and a 'Policy Rules' link. There are also '+' and 'x' icons for adding or removing rules.

System

Traffic Policies for Office

Default

Default Policy
The default policy defines the default action for traffic matching the pool. Exceptions to the default policy (i.e. using the opposite policy with respect to the default one) can be configured by adding rules below.

Pass

Policy Rules

L7 Protocol Rules
Configure here the L7/application protocol exception rules.

SSH Pass x IGMP Pass x MDNS Pass x sFlow Pass x AmazonAlexa Pass x SSDP Pass x ntop Pass x
Dropbox Pass x ICMP Pass x HTTP Pass x NetBIOS Pass x DNS Pass x Policy Rules + x

L7 Category Rules
Configure here the traffic category exception rules.

Malware Drop x Policy Rules + x

Host Rules
Configure here the hostname (e.g. used in TLS, DNS... protocols) exception rules.

emergingthreats.net Pass x s3.amazonaws.com Pass x centos.org Pass x snort.org Pass x github.com Pass x
abuse.ch Pass x ntop.org Pass x fedoraproject.org Pass x Policy Rules + x

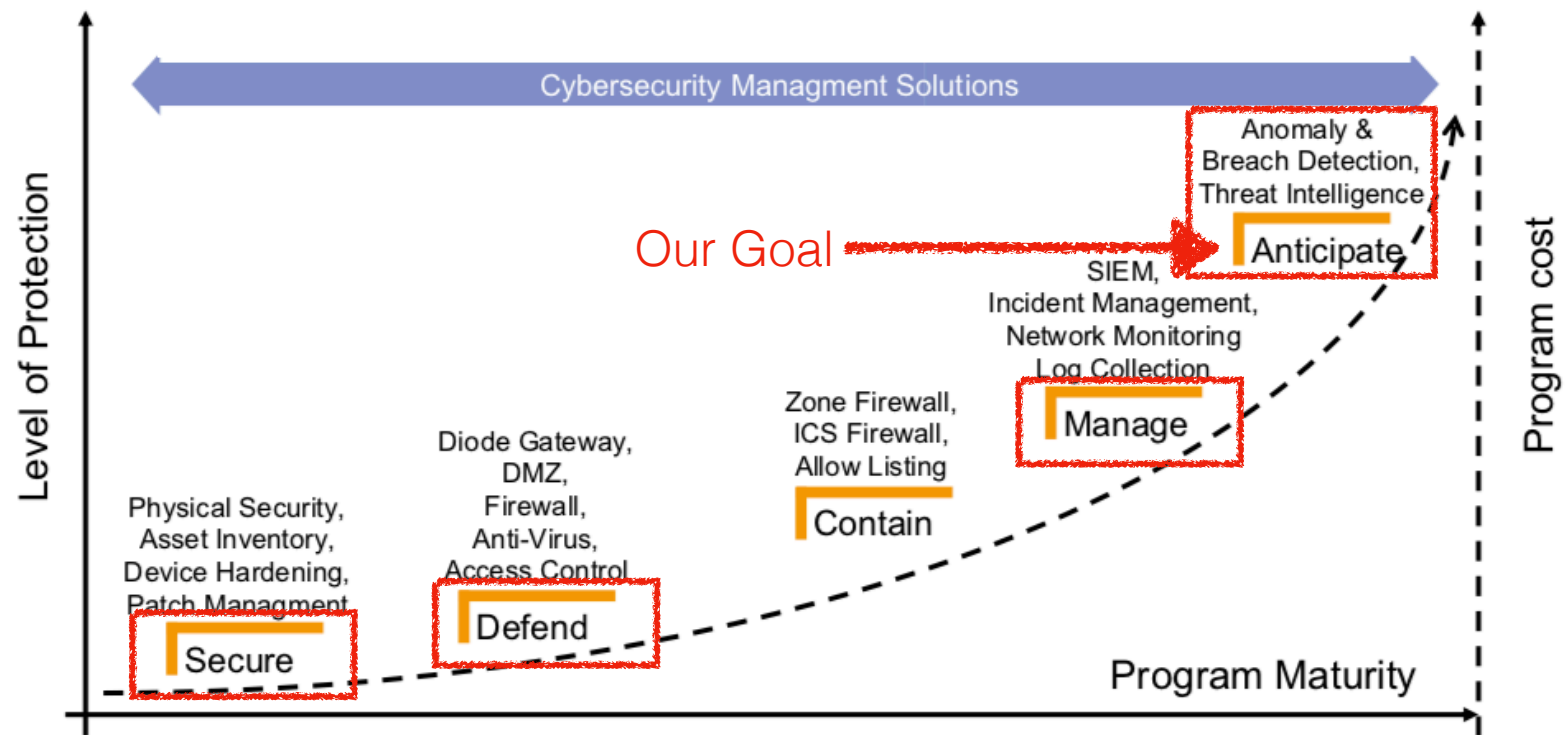
Countrv Rules

Using Score to Enforce Policies [7/7]

- The only exception to this policy is the Jailed Hosts pool:
 - Hosts added to this pool are blocked.
 - When a host is removed from this pool (after having been blocked), such host is moved back to the original pool (or the default pool).
- Whenever a policy is changed or a host is added/removed from this pool, ntopng informs all nProbes in IPS mode (yes, you can have more than one) automatically, with no user action whatsoever. All actions performed are logged, and “dry run” mode is available for simulating the actions before moving inline.

Part II: Ongoing Developments

2021 Monitoring Goals



Picture courtesy of [switch.ch](https://www.switch.ch)

How Can we Anticipate a Problem?

- Monitoring can show you when a problem is happening or (better) what are suspicious flows that can be an indication of a future problem.
- Can we do anything better than this? What if I could detect the user and application that generated a traffic flow?
- Goal: extend current monitoring capabilities with system analysis in order to report richer information and build new, more powerful checks.

Cybersecurity and Networking

- In a way, cybersecurity would not be that important without the Internet as networks propagate threats.
- Using DPI and traffic analysis techniques so far presented it is possible to have a great level of visibility and protection but...
- East-west traffic monitoring is not so simple and available techniques (e.g. sFlow) are sampled.
- Threats do their best to hide themselves: volumetric attacks are “nice” as they can be easily spotted.
- More packets, more ML and more checks are the only viable solution to this problem ?

nProbe Agent


- In 2018 we have released a Linux-only event based (i.e. packet-less) agent named nProbe Agent.

```
28/Apr/2019 23:46:29 [Netlink.cpp:1159] [Netlink] [counters] { "timestamp":  
"1556487989.626174", "ifName": "veth40297a6", "ifIndex": 21, "LOCAL_CONTAINER": { "DOCKER":  
{ "NAME": "tecmin-web3" } }, "ifInOctets": 32477, "ifInPackets": 328, "ifInErrors": 0,  
"ifInDrops": 0, "ifOutOctets": 13110951, "ifOutPackets": 40902, "EXPORTER_IPV4_ADDRESS":  
"x.x.x.x" }
```

```
9/Apr/2019 12:09:54 [EBPF.cpp:178] [eBPF] { "timestamp": "1556532594.175074", "LOCAL_PROCESS":  
{ "PID": 17932, "UID": 135, "GID": 145, "PROCESS_PATH": "/usr/bin/influxd" },  
"LOCAL_FATHER_PROCESS": { "PID": 1, "UID": 0, "GID": 0, "PROCESS_PATH": "/lib/systemd/  
systemd" }, "EVENT_TYPE": "ACCEPT", "IP_PROTOCOL_VERSION": 4, "PROTOCOL": 6, "L4_LOCAL_PORT":  
51176, "L4_REMOTE_PORT": 8086, "IPV4_LOCAL_ADDR": "127.0.0.1", "IPV4_REMOTE_ADDR": "127.0.0.1",  
"EXPORTER_IPV4_ADDRESS": "x.x.x.x" }
```

- The idea was to merge network with system visibility. It turned out to confuse people using nProbe. For this reason we have decided to revamp this idea:
 - Merging this code with nProbe and discontinue nProbe Agent.
 - Adding Windows support

Merging Network and System IEs [1/4]





[57640]	[Len 4]	%SRC_PROC_PID	Flow source process PID
[57641]	[Len 16]	%SRC_PROC_NAME	Flow source process name
[57897]	[Len 4]	%SRC_PROC_UID	Flow source process userId
[57844]	[Len 16]	%SRC_PROC_USER_NAME	Flow source process user name
[57845]	[Len 4]	%SRC_FATHER_PROC_PID	Flow source father process PID
[57846]	[Len 16]	%SRC_FATHER_PROC_NAME	Flow source father process name
[57847]	[Len 4]	%DST_PROC_PID	Flow dest process PID
[57848]	[Len 16]	%DST_PROC_NAME	Flow dest process name
[57898]	[Len 4]	%DST_PROC_UID	Flow dest process userId
[57849]	[Len 16]	%DST_PROC_USER_NAME	Flow dest process user name
[57850]	[Len 4]	%DST_FATHER_PROC_PID	Flow dest father process PID
[57851]	[Len 16]	%DST_FATHER_PROC_NAME	Flow dest father process name

Merging Network and System IEs [2/4]

Flow: 192.168.1.187:50837 ↔ 104.244.42.136:443 | Overview

Flow Peers [Client / Server]	192.168.1.187 :50837 ↔ 104.244.42.136 :443 [Twitter Inc.]	
Protocol / Application	TCP / TLS.Twitter (SocialNetwork)	
First / Last Seen	22/10/2021 16:16:06 [00:42 sec ago]	22/10/2021 16:16:06 [00:42 sec ago]
Total Traffic	Total: 93 Bytes —	
	Client → Server: 1 Pkts / 41 Bytes —	Client ← Server: 1 Pkts / 52 Bytes —
	<div><div>192.168.1.187:50837</div><div>104.244.42.136:443</div></div>	
DSCP / ECN [Client / Server]	Best Effort [CS0] / Disabled (0)	Best Effort [CS0] / Disabled (0)
CommunityId	1:/Z0Vhpynbl/B2TsC/c47Ei9ZAlg=	
Actual / Peak Throughput	0 bit/s — / 0 bit/s	
ASN [Client / Server]		13414
Flow Verdict	0	
Additional Flow Elements		
IPv4 address of the host were nProbe runs	192.168.1.187	
Client process name	C:\Program Files (x86)\Google\Chrome\Application\chrome.exe	
Total number of exported flows	373	

Merging Network and System IEs [3/4]



Merging Network and System IEs [4/4]

- Advantages

- Map traffic to processes/users: finally we know “who is doing what”.
- Detect unexpected processes making traffic.
- Simplified troubleshooting and incident analysis with contextual data.

- Limitations

- Still a passive tool: the collector has the knowledge.
- It is unable to detect “changes” but only “facts” (i.e. annotated flows with limited system metadata).

Towards a nProbe-based EDR


- What if nProbe could:
 - Detect changes in configuration invisible to the network.
 - Use process and user information to properly evaluate risks in communications.
 - Use contextual information (e.g. process) not just for enriching flow data but also for preventing threats from spreading in the network?
- What about a nProbe-based EDR (Endpoint Detection and Response) ?

Cybersecurity Simplified [1/2]

- Challenge: can we allow administrators to block threats before the problem shows up?
- Options: block traffic of applications that
 - Are not installed as package or that are started from non-standard locations (e.g. /tmp).
 - Have not been running previously.
 - Communicate with blacklisted IPs.
 - Have a periodicity and are not monitoring tools.
 - ...(cont).


Cybersecurity Simplified [2/2]

- Combining system visibility with network monitoring, enabled us to create an active probe able to block specific application traffic and that can very well fit with the zero-trust principle that is becoming increasingly popular.



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]

- It is possible to query redis for users who sent data out, and for each process (that transmitted/received data) run by each user.

```
$ redis-cli keys "process.*"  
1) "process.root"  
2) "process.www-data"  
3) "process.influxdb"  
4) "process._apt"  
5) "process.postgres"  
6) "process.avahi"  
7) "process.clickhouse"  
8) "process.chronograf"  
9) "process.deri"  
10) "process.grafana"
```

```
$ redis-cli hkeys "process.root"  
1) "/usr/sbin/NetworkManager"  
2) "/usr/lib/sm.bin/sendmail"  
3) "/usr/sbin/ntpdupdate"  
4) "/sbin/dhclient"  
5) "/usr/sbin/cups-browsed"  
6) "/snap/core/11606/usr/lib/snapd/snapd"  
7) "/home/deri/nprobe"  
8) "sendmail-mta"
```

- Is an unknown process allowed to do networking ?
Probably not.

Introducing System Visibility in nProbe [4/5]

- 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 not yet supported. We believe that osquery.io might be an option to consider in the future.

Introducing System Visibility in nProbe [5/5]

Flow: 192.168.1.178:56520 ↔ 192.168.1.187:22 Overview		
Flow Peers [Client / Server]	192.168.1.178 [R :56520 [28:37:37:00:6D:C8] ↔ 192.168.1.187 [R :22 [D8:CB:8A:E1:2D:2E]	
Protocol / Application	TCP / SSH (RemoteAccess) 🔒	
First / Last Seen	27/10/2021 16:56:35 [00:14 sec ago]	27/10/2021 16:56:36 [00:13 sec ago]
Total Traffic	Total: 684 Bytes —	
	Client → Server: 6 Pkts / 420 Bytes —	Client ← Server: 3 Pkts / 264 Bytes —
	<div><div>192.168.1.178:56520</div><div>192.168.1.187:22</div></div>	
DSCP [ECN] [Client / Server]	Immediate [AF21] / Disabled (0)	Unknown [4] / Disabled (0)
RTT Time Breakdown	<div><div>0.165 ms (client)</div><div>0.182 ms (server)</div></div>	
Max (Estimated) TCP Throughput	Client → Server: 94.43 Mbit/s	Client ← Server: 11.55 Mbit/s
TCP Flags	Client → Server: A P	Client ← Server: A P
	Flow is active, however, the beginning of the flow has not been seen and peer roles (client/server) might be inaccurate	
Total Flow Score / Score Category Breakdown	10	<div>Network</div>
Issues	Description	Actions
	Remote Access [Score: 10] ⚠️	<div><div>🚫</div><div>⚙️</div><div>⚠️</div></div>
CommunityId	1:6kPbNQwvDTagswGSa8ETWbGyegA=	
Actual / Peak Throughput	5.47 kbit/s — / 5.47 kbit/s	<div></div>
Flow Verdict	0	
Additional Flow Elements		
Flow Exporter IPv4 Address	192.168.1.187	
nDPI Flow Risk Score	0	
Client Process	/usr/sbin/sshd	
Client Process Package	openssh-server	

Further Visibility: Server Side [1/3]

- As said before, a good strategy for detecting issues/reconfigurations/malware is to track changes.
- When a malware speaks with remote peers, nProbe can detect the flow and report contextual information (process and package name).
- What if the malware isn't making any traffic (so it's in essence invisible to flows) but it's ready to accept connections from applications? Or if the traffic is so little that hides itself in background noise?

Further Visibility: Server Side [2/3]

- nProbe has been enhanced with local host port monitoring for:
 - Binding a port with an application and a package.
 - Detecting changes in port allocation: a new port is open, an existing port is closed, or a different process is listening to an existing open port.
 - Reporting this information to flow collectors for increased visibility.
- This feature is implemented on both Windows and Linux nProbe versions.

Further Visibility: Server Side [3/3]

```
{
  "ip-addresses": ["10.3.240.28", "192.168.1.187"],
  "listening-ports": {
    "tcp4": [{
      "port": 22,
      "proc": "/usr/sbin/sshd",
      "pkg": "openssh-server"
    }, {
      "port": 53,
      "proc": "/usr/sbin/dnsmasq",
      "pkg": "dnsmasq-base"
    }, {
      "port": 1234,
      "proc": "/home/deri/nProbe/nprobe",
      "pkg": "" ← No Package !
    }
  ],
  "tcp6": [{
    "port": 9000,
    "proc": "/usr/bin/clickhouse",
    "pkg": "clickhouse-common-static"
  ]
}
```

...

Exporting System Information

- Process information can be combined with DPI and flow risks to determine the flow “verdict”.

```
# cat /tmp/2021/09/22/22/49.flows
IPV4_SRC_ADDR|IPV4_DST_ADDR|INPUT_SNMP|OUTPUT_SNMP|IN_PKTS|IN_BYTES|FIRST_SWITCHED|LAST_SWITCHED|L4_SRC_PORT|
L4_DST_PORT|TCP_FLAGS|PROTOCOL|SRC_PROC_NAME|SRC_PROC_PID|DST_PROC_NAME|DST_PROC_PID|FLOW_VERDICT
192.168.1.187|192.168.1.178|0|0|17|6564|1632343764|1632343765|22|56218|24|6|0|0|usr/sbin/sshd|2910|0
192.168.1.178|192.168.1.187|0|0|17|884|1632343764|1632343765|56218|22|16|6|0|0|usr/sbin/sshd|2910|0|0
192.168.1.178|192.168.1.187|0|0|9|612|1632343767|1632343768|49372|22|24|6|0|0|usr/sbin/sshd|2910|0|0
192.168.1.187|192.168.1.178|0|0|5|504|1632343767|1632343768|22|49372|24|6|0|0|usr/sbin/sshd|2910|0
192.168.1.187|192.168.1.178|0|0|11|3648|1632343767|1632343768|22|56218|24|6|0|0|usr/sbin/sshd|2910|0
192.168.1.178|192.168.1.187|0|0|11|572|1632343767|1632343768|56218|22|16|6|0|0|usr/sbin/sshd|2910|0|0
192.168.1.187|192.168.1.1|0|0|2|116|1632343768|1632343768|44199|53|0|17|0|0|usr/bin/traceroute.db|4909|2
192.168.1.1|192.168.1.187|0|0|1|106|1632343768|1632343768|53|44199|0|17|0|0|usr/bin/traceroute.db|4909|0|2 ← Drop
192.168.1.187|192.168.1.178|0|0|9|3264|1632343771|1632343771|22|56218|24|6|0|0|usr/sbin/sshd|2910|0
192.168.1.178|192.168.1.187|0|0|9|468|1632343771|1632343771|56218|22|16|6|0|0|usr/sbin/sshd|2910|0|0
192.168.1.178|192.168.1.187|0|0|4|244|1632343772|1632343772|49372|22|24|6|0|0|usr/sbin/sshd|2910|0|0
192.168.1.187|192.168.1.178|0|0|3|296|1632343772|1632343772|22|49372|24|6|0|0|usr/sbin/sshd|2910|0
```

- **2** means drop as traceroute was either unknown during learning phase, or not part of an installed package (this culprit can be solved if SRC_PROC_PACKAGE_NAME is also exported).

Enforcement vs Monitoring [1/2]

- nProbe can both enforce traffic policies (i.e. pass/drop) or passively monitor traffic.
- The difference is just on how the tool is started:
 - Monitoring
 - Capture traffic from an interface.
 - Enforcement
 - nProbe is started on top of netfilter (Linux firewall architecture) for blocking traffic if necessary.

Enforcement vs Monitoring [2/2]

- Passive Monitoring

```
nprobe -i enp5s0 -T "%IPV4_SRC_ADDR %IPV4_DST_ADDR %INPUT_SNMP %OUTPUT_SNMP  
%IN_PKTS %IN_BYTES %FIRST_SWITCHED %LAST_SWITCHED %L4_SRC_PORT %L4_DST_PORT  
%TCP_FLAGS %PROTOCOL %SRC_PROC_NAME %DST_PROC_NAME %FLOW_VERDICT" --redis  
localhost --process-learning-duration 86400:0
```

- Enforcement

```
nprobe -i nf:0 -T "%IPV4_SRC_ADDR %IPV4_DST_ADDR %INPUT_SNMP %OUTPUT_SNMP %IN_PKTS  
%IN_BYTES %FIRST_SWITCHED %LAST_SWITCHED %L4_SRC_PORT %L4_DST_PORT %TCP_FLAGS  
%PROTOCOL %SRC_PROC_NAME %DST_PROC_NAME %FLOW_VERDICT" --redis localhost --  
process-learning-duration 86400:0
```

- This nProbe pre-release is currently available for Ubuntu 18.04 and 20.04.
- Windows version of nProbe is (so far) monitoring only.
- Final release is expected in December/January timeframe.
- Note: all items discussed so far are container friendly.

Open Discussion



- <https://blog.ntop.org>
- <https://github.com/ntop/>
- <https://www.ntop.org/community/>