Welcome to ntopConf 2022

Luca Deri <deri@ntop.org>
@lucaderi

Introduction

Who am I

- ntop founder (http://www.ntop.org): company that develops open-source network security and visibility tools.
- Author of various open source software tools.
- Lecturer at the CS Dept, University of Pisa.
- ntop tools are used by leading educational institutions, trading firms, cybersecurity companies.

20+ Years of ntop

- Private company focusing on network traffic monitoring, security and high-speed networking.
- In 1998 we have released the original ntop, an open source web-based network monitoring application.
- Today we develop in-house various products both open source (https://github.com/ntop/) and proprietary.
- Thanks to open-source and to our policy to give free software to education, ntop is a well known brand in this market and many universities use our tools.
- In 2023 we'll celebrate our 25th anniversary

Why ntop?

- Software developed from the ground up: kernel drivers, application, libraries. Everything is under our control.
- No external dependencies: price and features won't be a surprise.
- Two decades in business: we plan to stay around.
- Vendor neutral: we want to offer you what is the best available, with no hidden vendor dependencies.
- Multi-platform support: Linux, Windows, MacOS, FreeBSD.

Us and the OpenSource

- We have always been opensource in most of our products.
- We have created a non-open source product line for both funding us (without relying on training or donations) and entering business environments with our tools.
- We authored other leading open source tools including Wireshark (Traffic Analyzer) and Suricata (IDS/IPS).
- We produce open-source components used by many other opens source products.
- We have always donated our tools to NGOs, no-profit, research institutions, universities and education.

International Sales Presence

Americas

- FirstLight [USA/Canada]

Europe

- ~~ByteLabs [Russia/CIS]~~
- Gravitate [Germany]
- Hosting Solutions [Europe]
- Info-Stor [UK, Nordics]
- Loops Cloud Computing [Catalunya]
- Lugos [France]
- Miniserver [Europe]
- quattroSEC [Austria]
- Teknoservice [Europe]
- verXo – Security solutions [Europe]
- Vunkers [Spain]
- Würth-Phoenix [Europe]

Asia / APAC / Middle East

- Assured Network Solutions [Australia/New Zealand]
- Info-Stor [India/Pakistan and neighbouring countries]
- Linksoft [Taiwan]
- Hongke Technology [China/Taiwan]
- Jupiter Technology Corp. [Japan]
- npacket [Korea]
- Softense [Israel]

ntop R&D

- Italy

International Sales

- Switzerland

Commercial Products

eShop

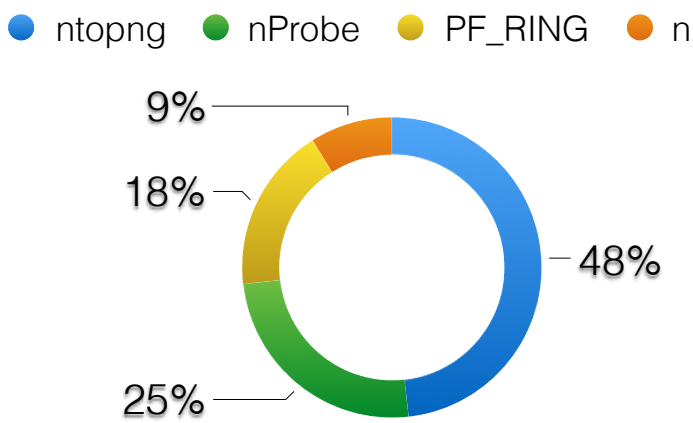

OEM
Direct Sales

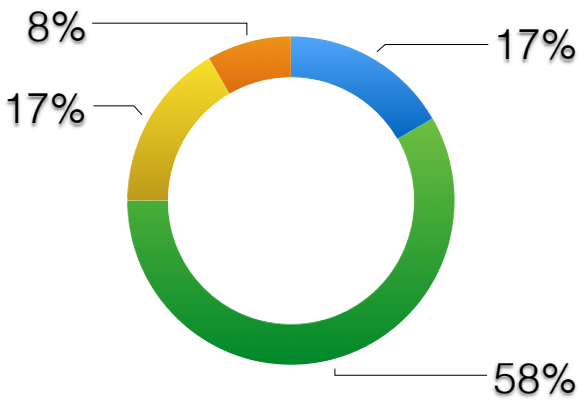

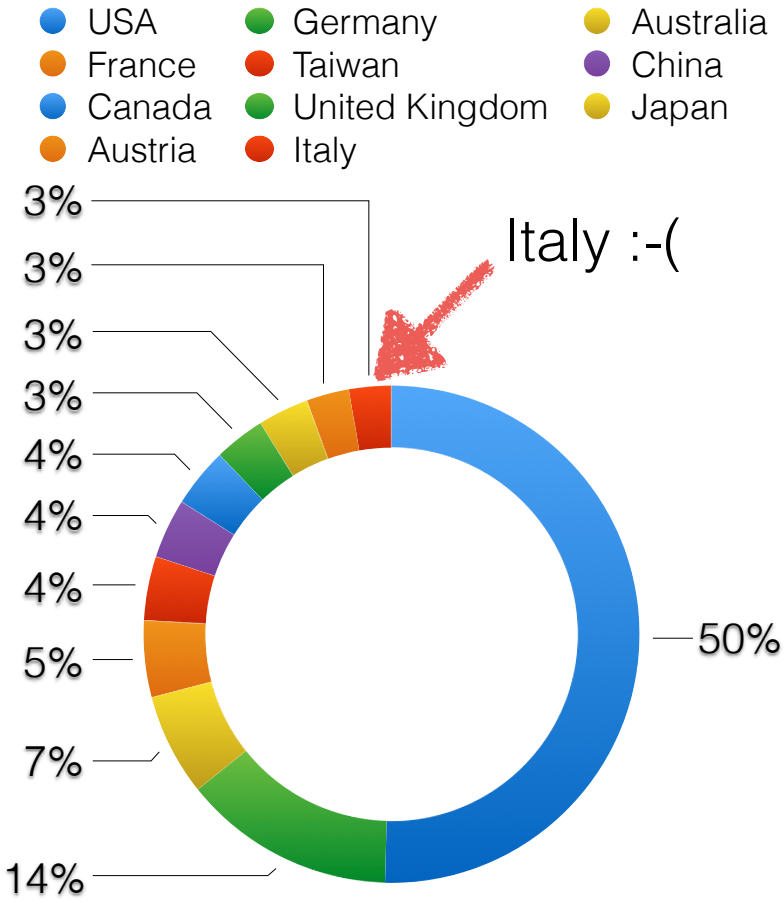

From ntopConf 2019 to 2022+

ntop 2019 in Retrospective

- Focus on traffic visibility
  - nDPI for application protocol analysis and traffic accounting
  - ntopng 3.9: high-resolution counters, long-term timeseries storage and flow-indexing.
  - Information DataSource: integration with Suricata, SecurityOnion, third-party vendors (Fortinet, SonicWall, Sophos, and others) to provide monitoring data and receive signature-based alerts.
  - Early experiments on Linux systems visibility via eBPF.

ntop Trends in 2019-22

- Focus on cybersecurity traffic analysis.
- Delivered professional training as many professional users asked it.
- Scalability and performance bump 10/40 -> 100: ntopng can now monitor in realtime 100 Gbit or tenth of 10/40 Gbit networks leveraging on improved nProbe and flow-transfer protocols (encryption + data compression).
- IoT and ICS/SCADA monitoring: better understanding of “home” IoT (visibility and threats), analysis of industrial protocols.
- Behavioural/statistical traffic analysis: not just show what, but interpret data (i.e. visibility is no longer enough).

Traffic Signatures and ML [1/2]

- Traditional security tools, 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.

Traffic Signatures and ML [2/2]

- Increasing adoption of encryption technologies, DPI 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.
- At ntop we believe that lightweight statistical algorithms can produce the same results of heavier ML-based solutions.
- The above assumption holds if we're able to analyse encrypted traffic (ETA) that on Internet links exceeds 80% of the total.

From Academy to Practice

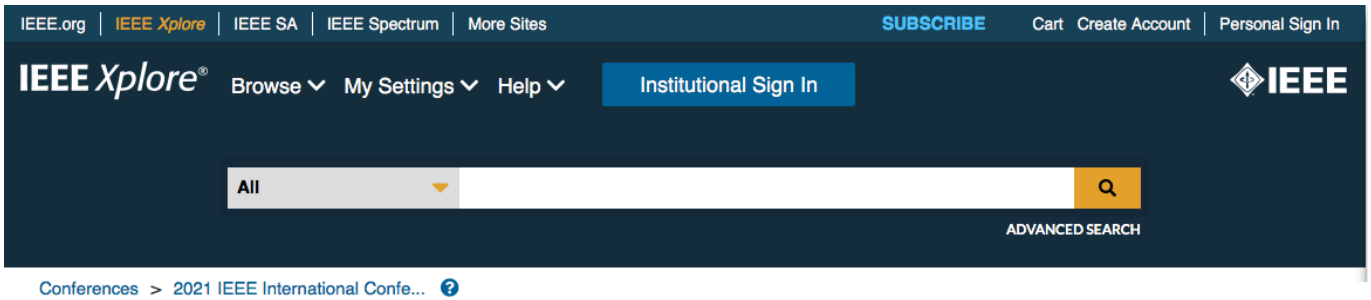

2021

2022 (To Appear)

Using CyberScore for Network Traffic Monitoring

Luca Deri
ntop
Pisa, Italy
deri@ntop.org

Alfredo Cardigliano
ntop
Pisa, Italy
cardigliano@ntop.org

***Abstract*—The growing number of cybersecurity incidents and the always increasing complexity of cybersecurity attacks is forcing the industry and the research community to develop robust and effective methods to detect and respond to network attacks. Many tools are either built upon a large number of rules and signatures which only large third-party vendors can afford to create and maintain, or are based on complex artificial intelligence engines which, in most cases, still require personalization and fine-tuning using costly service contracts offered by the vendors.**

This paper introduces an open-source network traffic monitoring system based on the concept of cyberscore, a numerical value that represents how a network activity is considered relevant for spotting cybersecurity-related events. We describe how this technique has been applied in real-life networks and present the result of this evaluation.

***Index Terms*—Security Score, Deep Packet Inspection, Network Intrusion Detection, Traffic Measurement, Open-Source.**

I. Introduction and Motivation

for developing a novel method able to combine behavioural traffic analysis and encrypted network traffic analysis via DPI (Deep Packet Inspection), which is based on statistical methods and thus light and simple to operate in comparison to AI-based systems [6]. ntopng [4] is an open-source network traffic monitoring system developed by the authors, that leveraging on nDPI [5], an open-source DPI toolkit, can inspect and analyze network traffic, detect security issues and track host activities. In ntopng we have implemented the concept of *cyberscore* that is a numerical relevance indicator assigned to every observed network activity: the higher the value is, the higher the severity of this activity. Every network flow is inspected using nDPI, which can detect and report a set of flow risks whenever an unexpected issue, such as an expired TLS certificate, a DGA/Punycode/IDN domain name or credentials transfer in clear text, is detected. To date, nDPI support 45 different flow risks classification for both clear-text and

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.
- Each risk has associated a client/server score that is a numerical value used to indicate how severe is the reported issue.

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
- Possible Exploit Detected
- TLS Certificate Close to Expire
- Punycode/IDN Domain
- Error Code Detected
- Crawler/Bot Detected
- Anonymous Subscriber

Legenda: Clear Text Only, Encrypted/Plain Text, Encrypted Only

From Flow Risk To Score

TCP 192.168.149.129:43535 <-> 51.83.238.219:80 [proto: 91.252/**TLS.AnyDesk**][cat: RemoteAccess/12][2942 pkts/175103 bytes <-> 4001 pkts/2618640 bytes][Goodput ratio: 9/92][55.97 sec][bytes ratio: -0.875 (Download)][IAT c2s/s2c min/avg/max/stddev: 0/0 19/14 7028/7028 153/126][Pkt Len c2s/s2c min/avg/max/stddev: 54/60 60/654 1514/1514 50/618][Risk: ** Known protocol on non standard port **** TLS (probably) not carrying HTTPS **** SNI TLS extension was missing **** Desktop/File Sharing Session **][**Risk Score: 80**][TLSv1.2][JA3C: 201999283915cc31cee6b15472ef3332][JA3S: 107030a763c7224285717ff1569a17f3][Issuer: CN=AnyNet Root CA, O=philandro Software GmbH, C=DE][Subject: C=DE, O=philandro Software GmbH, CN=AnyNet Relay][Certificate SHA-1: 9E:08:D2:58:A9:02:CD:4F:E2:4A:26:B8:48:5C:43:0B:81:29:99:E3][Firefox][Validity: 2018-11-18 02:14:23 – 2028-11-15 02:14:23][Cipher: TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384][Plen Bins: 0,7,17,2,1,2,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,2,0,3,0,0,33,0,0,0,0,29,0,0]

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 [1/2]

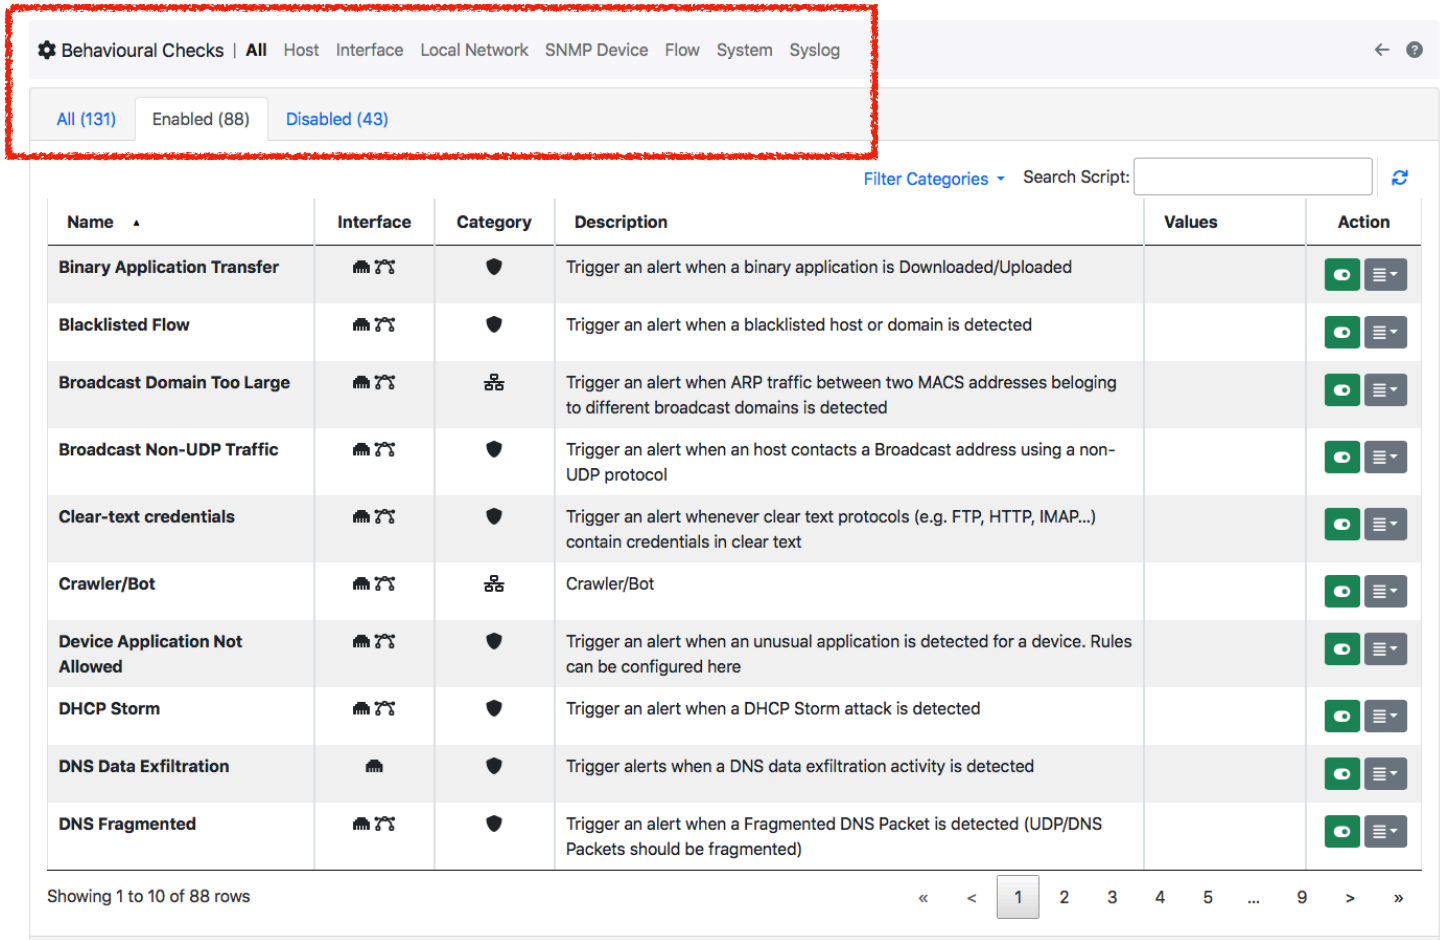

Behavioural Checks | All Host Interface Local Network SNMP Device Flow System Syslog

All (131) Enabled (88) Disabled (43)

Filter Categories Search Script:

Name	Interface	Category	Description	Values	Action
Binary Application Transfer			Trigger an alert when a binary application is Downloaded/Uploaded		
Blacklisted Flow			Trigger an alert when a blacklisted host or domain is detected		
Broadcast Domain Too Large			Trigger an alert when ARP traffic between two MACS addresses beloging to different broadcast domains is detected		
Broadcast Non-UDP Traffic			Trigger an alert when an host contacts a Broadcast address using a non-UDP protocol		
Clear-text credentials			Trigger an alert whenever clear text protocols (e.g. FTP, HTTP, IMAP...) contain credentials in clear text		
Crawler/Bot			Crawler/Bot		
Device Application Not Allowed			Trigger an alert when an unusual application is detected for a device. Rules can be configured here		
DHCP Storm			Trigger an alert when a DHCP Storm attack is detected		
DNS Data Exfiltration			Trigger alerts when a DNS data exfiltration activity is detected		
DNS Fragmented			Trigger an alert when a Fragmented DNS Packet is detected (UDP/DNS Packets should be fragmented)		

Showing 1 to 10 of 88 rows « < 1 2 3 4 5 ... 9 > »

Consolidating Score [2/2]

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

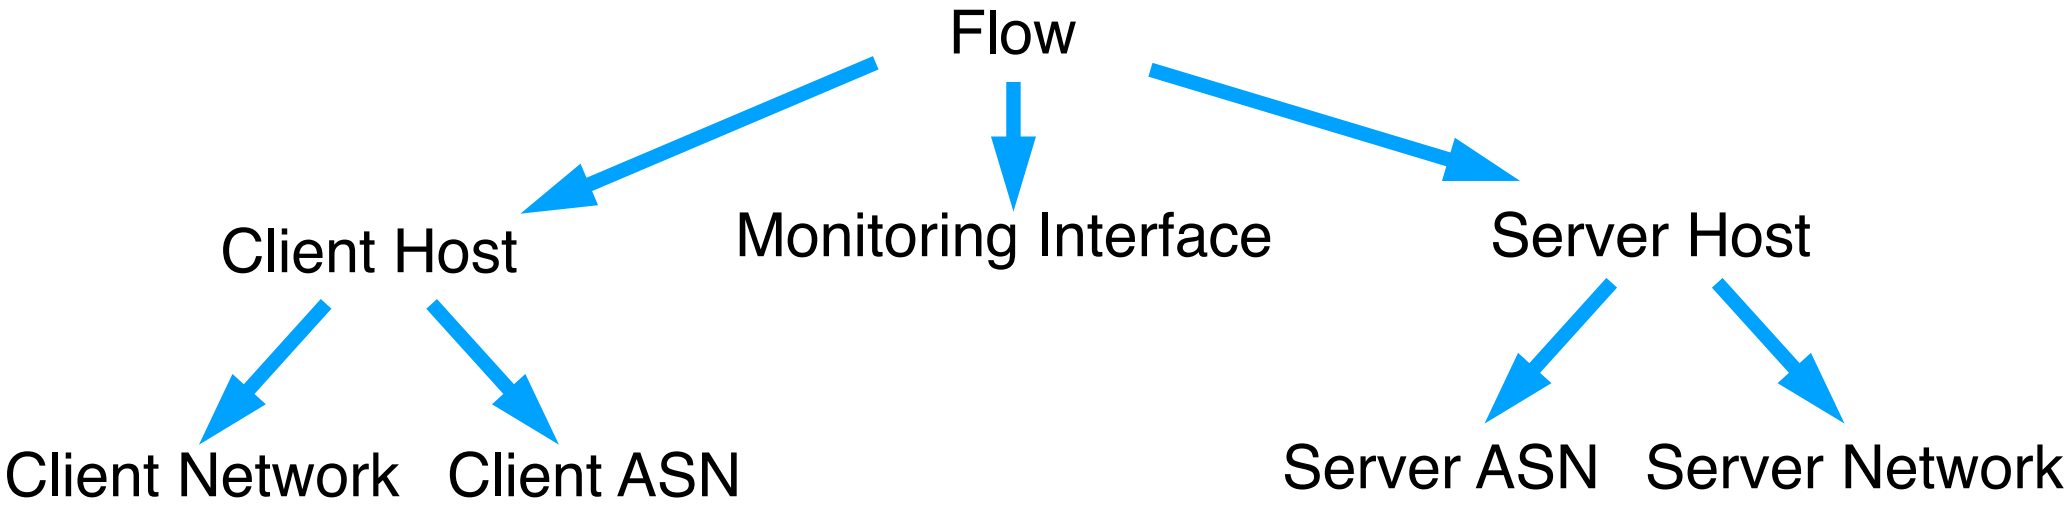

Visualising Cybersecurity: Bubbles

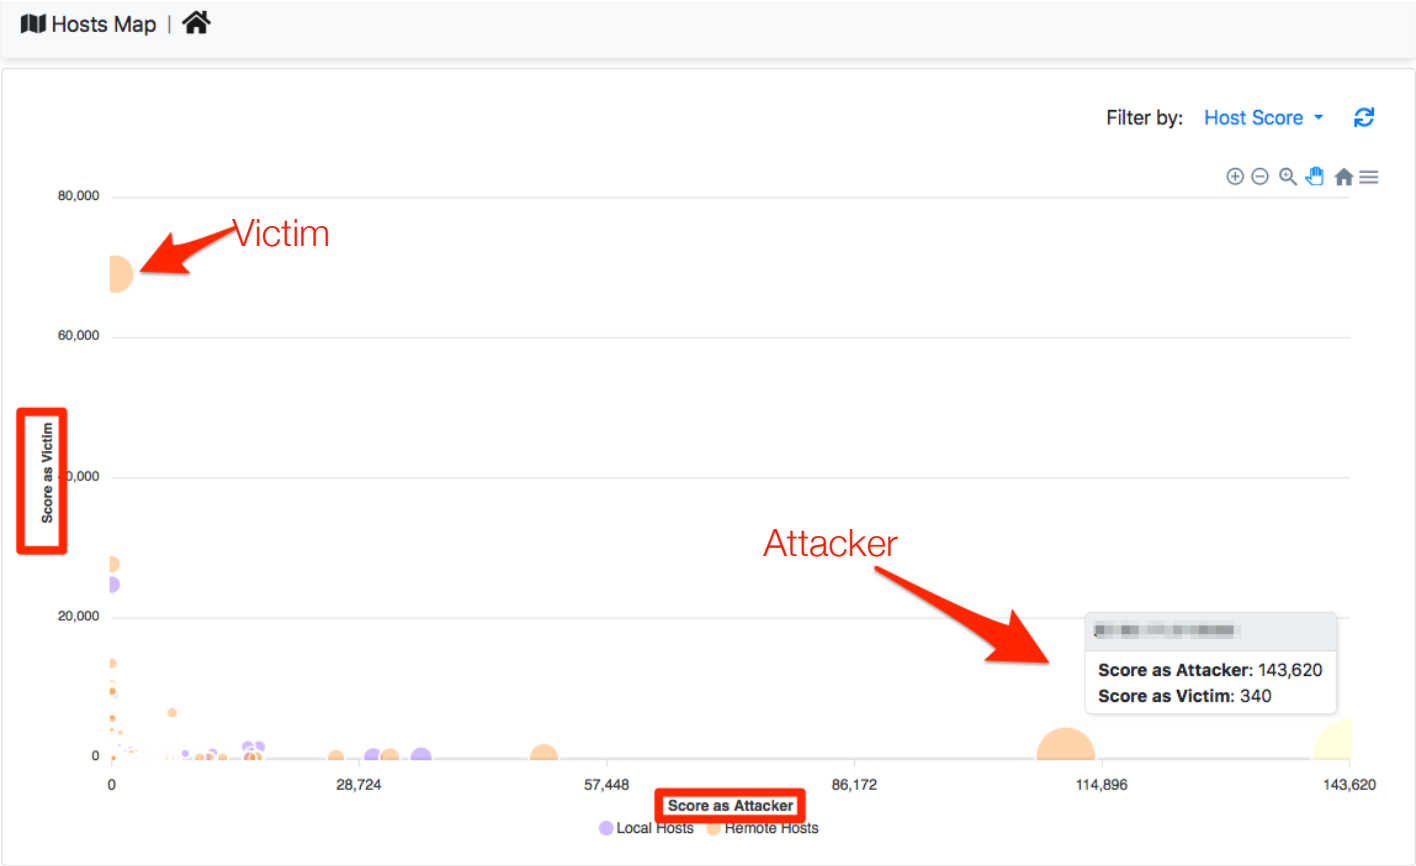

Score-based Behaviour Analysis [1/3]

- 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 [2/3]

- 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.

Score-based Behaviour Analysis [3/3]

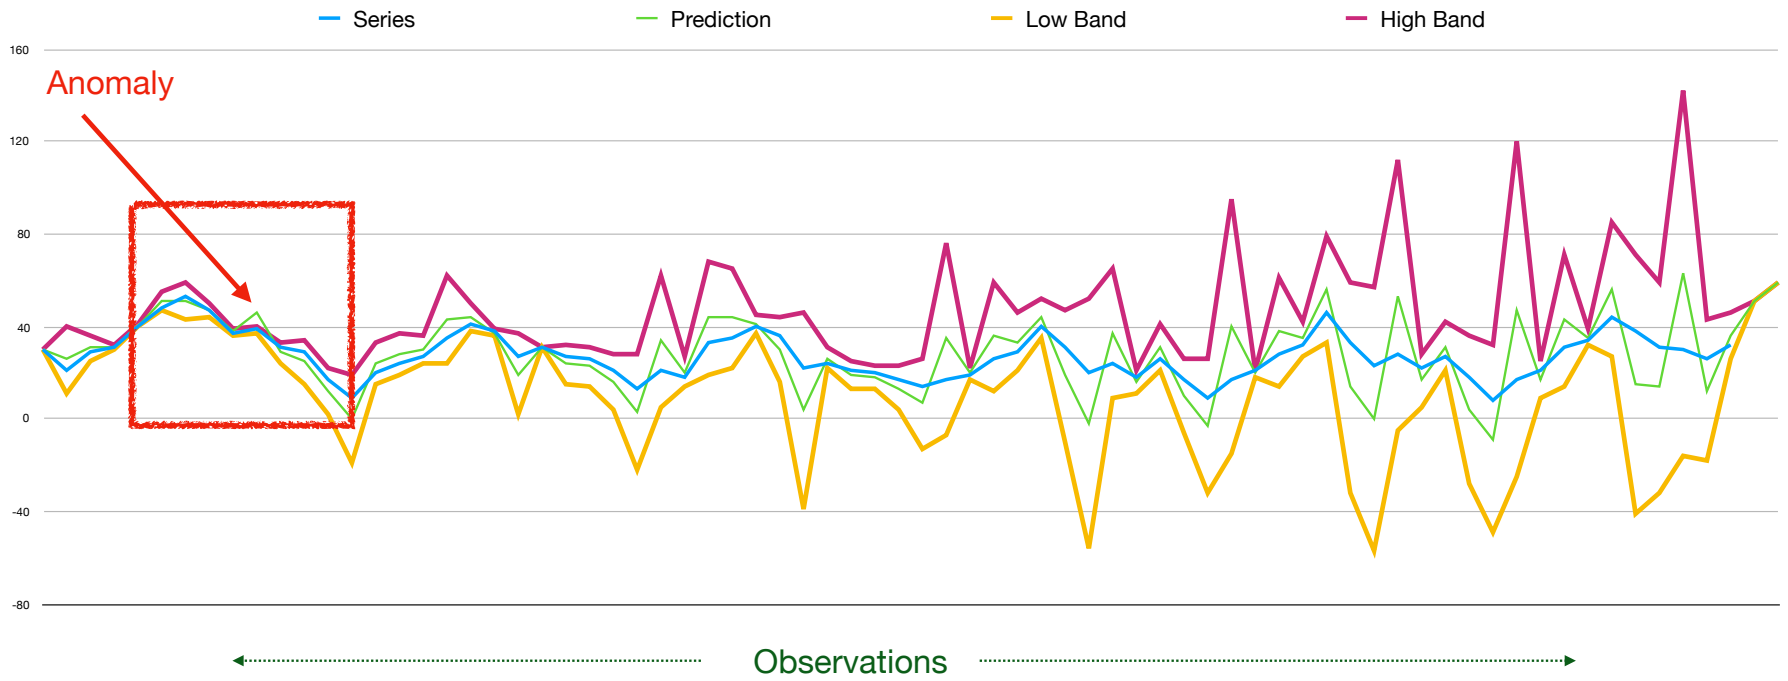

Using Learning in Behaviour Analysis

- In real life, our past experience determines our future behaviour.
- If we want to anticipate problems, we need to spot changes in behaviour at various levels:
  - Lateral Movements: talk with peers in the network.
  - Periodic Service Usage: find recurrent usage patterns (beaconing).
  - Application Popularity: spot applications used very seldom and bind them to network traffic.

Cybersecurity Simplified

- 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 (learning).
  - Communicate with blacklisted IPs (anticipate).
  - Have a periodicity and are not monitoring tools.

Merging Network with System Visibility

- Application visibility is important as it allows us to better characterise traffic and enrich it with process/user/container/pod information.
- In container analysis we can even rely just on system events without capturing traffic (costly and often not necessary on such dynamic environments).
- In order to achieve all this, we need to deploy nProbe on the monitored host(s) in order to report traffic and system information and deliver them as standard, albeit richer, IPFIX/NetFlow flows.

Using nProbe as a Host-based EDR

- What if we 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?
- In essence turn nProbe from a pure passive component into an active (optional) component able to provide richer data and act as an enforcer.

IoT/ICS: Visibility and Cybersecurity

- We have added initial ICS support in our tools (IEC104) in order to detect behavioural changes in industrial systems.
- Experience has shown that:
  - Behaviour analysis is key to spot issues.
  - Changes in healthy ICS networks must be promptly detected as they can hide threats or faults.

ntop 2022 Roadmap (Open Discussion)

- From Software to Service
  - Currently we support multi-tenancy that enable ISPs/Providers to deliver traffic analysis to their customers using a single instance.
  - We need to enable service providers to implement services based on ntop tools hence
    - Enable NOC/SOC companies to use ntop-provided data for visibility and security.
    - Hierarchical monitoring/management of large networks from a single vantage point.
- More Visibility, Cybersecurity, ICS/Scada

ntop 2022 Roadmap (Open Discussion)

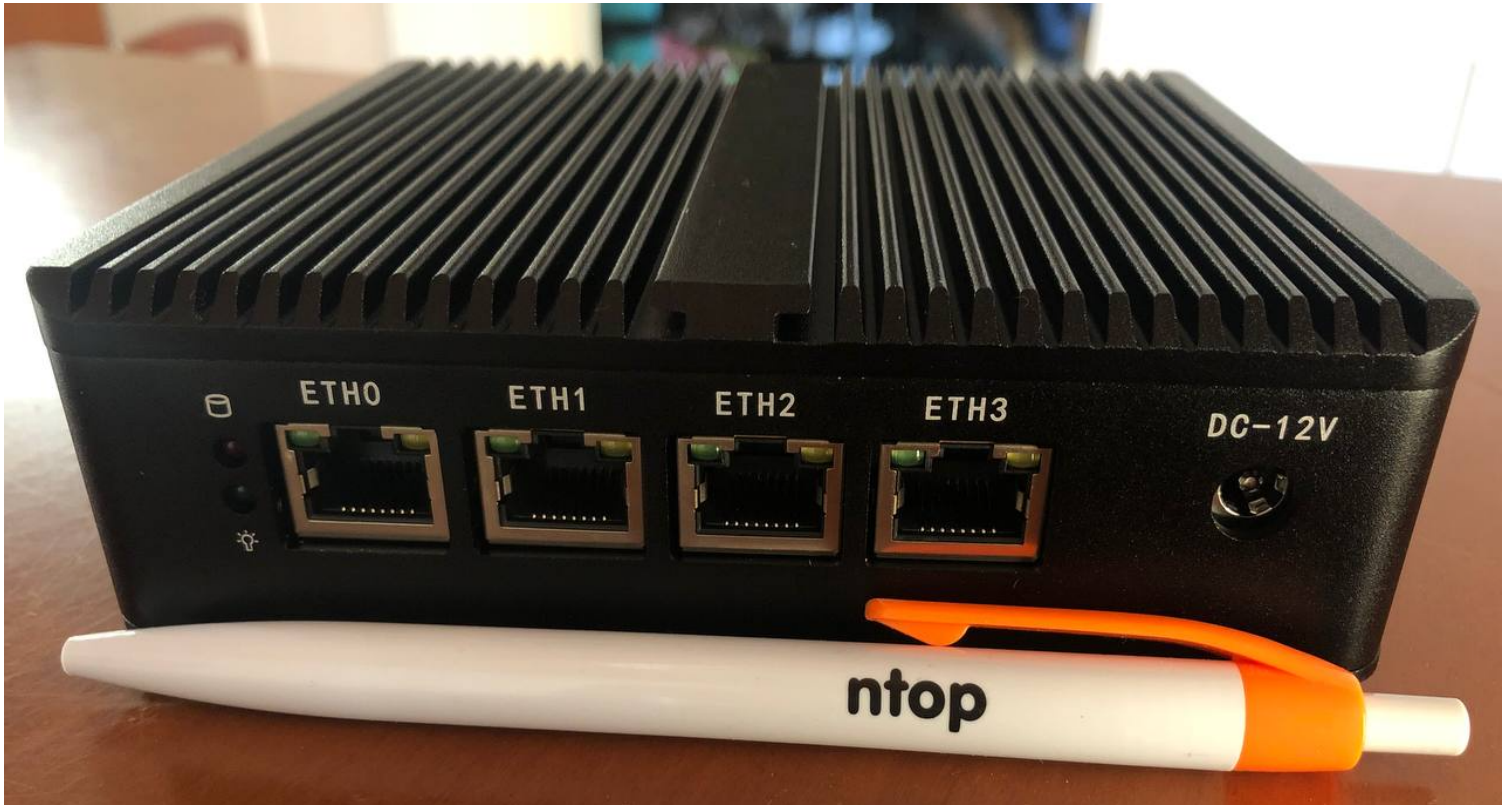

ntop 2022 Roadmap Announcements

- Explore the combination of ntop passive scanning with OpenVAS active Scanning. We are already in contact with Greenbone Networks, the makers of OpenVAS, to deliver by end 2022 an European, open source passive and active cybersecurity monitoring tool.

- Integration of ntop tools in Catchpoint products to combine network traffic visibility with active user experience monitoring.

Final Remarks

- Cybersecurity and system visibility are the driving forces of our current developments.
- We have consolidated our tools coming them from 10/40 to 100 Gbit traffic analysis, still using commodity hardware.
- We have open our tools to both act as source of data for third party tools (e.g. Elastic), and accept ingress feeds (e.g. Suricata or third-party vendors) for enriching the network view.
- We have contributed to open source, while offering tools with commercial support able to serve the needs of modern enterprises.
- Roadmap discussion at the end of ntopConf 2022.

Joining ntop

jobs@ntop.org