tnc23
DIGITAL GENERATIONS
TIRANA, ALBANIA | 5-9 JUNE 2023

Using AI/ML for Network-optimized DDoS Mitigation

Based on a DDoS botnet study in 2022-2023

Jérôme Meyer (Nokia)

Tirana

7 June 2023

GÉANT

Co-funded by the European Union

Then and now…
#1: Botnets have taken over the (DDoS) world

2002 - 2022

- Majority of DDoS is spoofed (IP header modification, IPHM)
- Originates from ~50 EU / AP hosting providers
- Abuses misconfigured NTP / DNS servers

2023

- We looked at (thousands of) attacks in 2022-2023

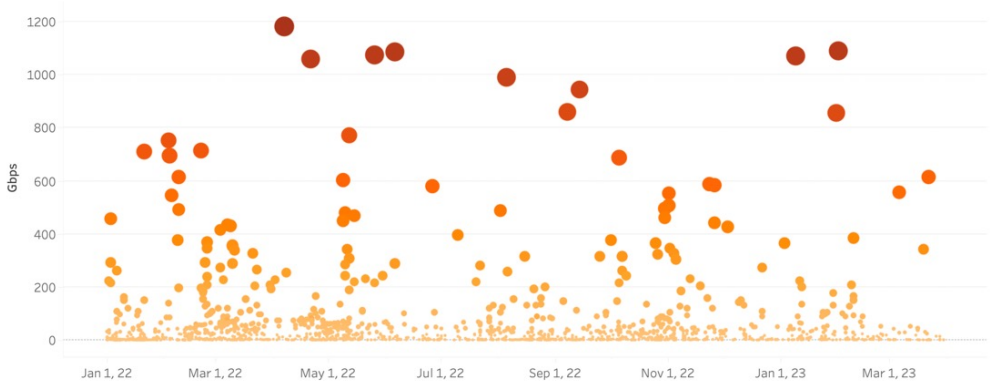

2022 - 2023 Nokia Botnet Study

Deepfield Secure Genome

Crawl every IPv4 and active IPv6 from datacenters around the world for known botnet CVE, IoT devices, servers and services

Nokia Deepfield Defender

Commercial DDoS security solution deployed in major ISP and Cloud providers around the world. Real-time telemetry from backbone routers and mitigation devices / scrubbers

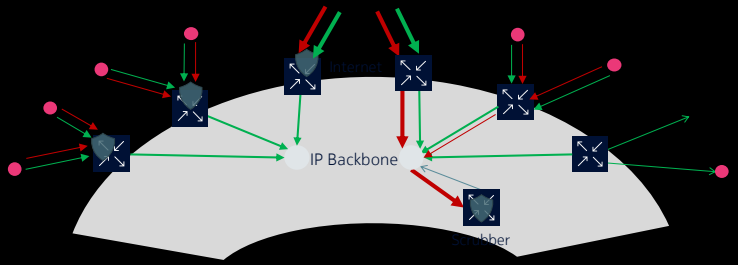

2022 - 2023 Nokia Botnet Study

Global DDoS Threat Alliance

Nokia collaboration with CSP / Nokia customers around the world sharing near real-time data on botnets, worms, DDoS attacks and other critical security threats

Then and now…
#1: Botnets have taken over the (DDoS) world

Today (now):

- Botnets generate most of all DDoS bytes
- Botnets represent **90% of complex attacks**
- **Botnets circumvent traditional anti-DDoS systems**

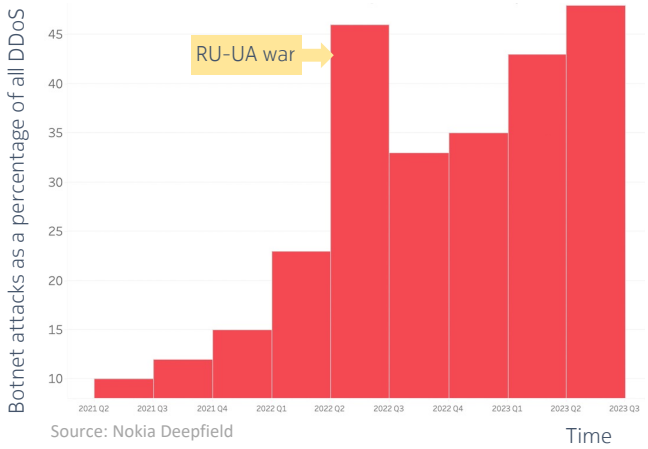

The graph shows Nokia data about botnet-originated DDoS traffic as a percentage of all attack traffic over the last year.
Data source: GDTA-participating service and cloud providers globally, using the Nokia commercial DDoS defense solution.

DDoS bots: What are they, and where do they live?
#2: The threats are growing exponentially, too

Where are these bots?

- Enterprises: IoT and cloud are now everywhere
  - Surveillance / Digital Video Recorders / Network Video Recorders
  - Point of Sale, Heating/Ventilation/AC, remote monitoring and data collection (water meters, parking meters)
  - Medical imaging

What are these bots?

- Most bot devices are compromised CPE (e.g., Mikrotik router), followed by 30-40 brands of DVR
- Botnets tend to attack in “packs” (similar devices and topologies)
- **Cloud is *not* the largest source (by number of devices), but one of fastest growing in terms of bandwidth (bps) and packet intensity (pps) capacity**

It may/will get worse:

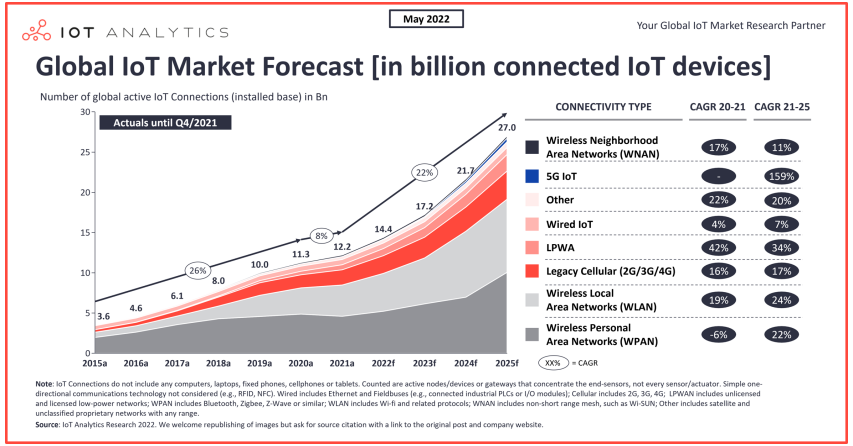

Source: https://iot-analytics.com/number-connected-iot-devices

- 99% of enterprise IoT and properly patched, firewalled and secure,
- but….
- **1% of many billion devices is significant.**

How big is the problem?
#3: Thousands of botnets, hundreds of thousands of bot devices

Today, based on Nokia data (and others), botnet DDoS represents:

- **500k – 1M** active IoT hosts
- **50 – 100 Tbps** aggregate capacity
- **1 – 2 Tbps** peaks

How many bots and botnets?

- **Majority attacks < 5,000 devices** and deliver effective attacks on many servers/applications
- There are large networks with **> 60k devices**
- Geo-political attacks included previously unknown botnet devices

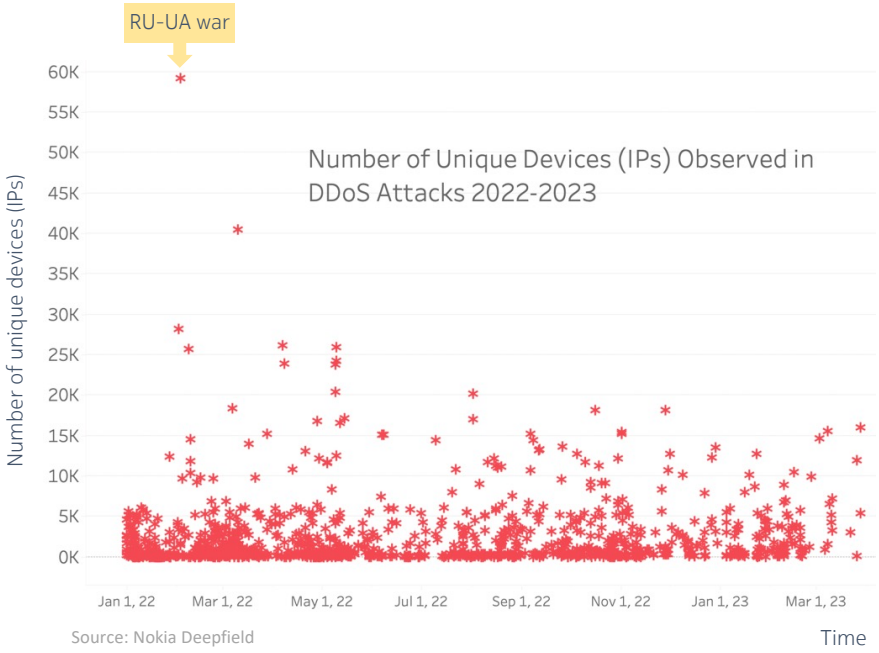

Some facts:
#4: Yet we're still in the early stages of botnet-driven DDoS impact

Last 20 years of Internet history

- Most (consumer/SMB) access via cable/DSL
- Asymmetric access 90 Mbps/10 Mbps (down/up)

Botnet threat is still limited

- Botnet bps matches industry averages
- 70% of all botnets < 50 Mbps today

But…

So far, **botnets are limited by today's upstream bandwidth** — while the race to gigabit speeds and symmetrical bandwidth is already well underway.

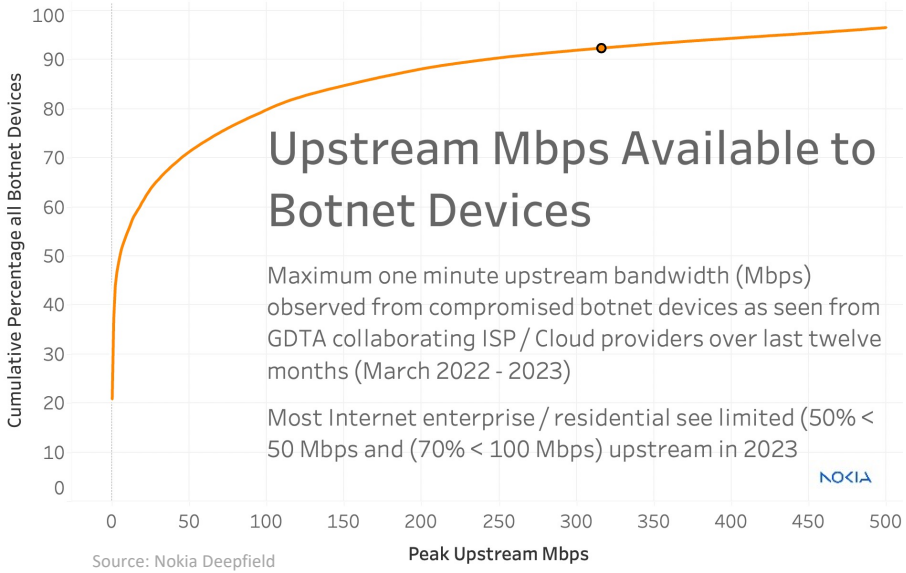

Why botnet attacks are such a problem?
"The call is coming from inside the house"

Traditional ISP / CSP security model assumed:

- Protect external edges of the network from inbound attacks,
  - Especially problematic in Eastern EU / Asian countries
- Protect against spoofed or amplified traffic
  - Active countermeasures (e.g., SYN cookie, HTTP redirect)
  - Shaping DNS, NTP, LDAP

The reality in 2023:

- Majority of botnet problem is North America / Europe
- **Largest threat for many ISP is from their own customers**

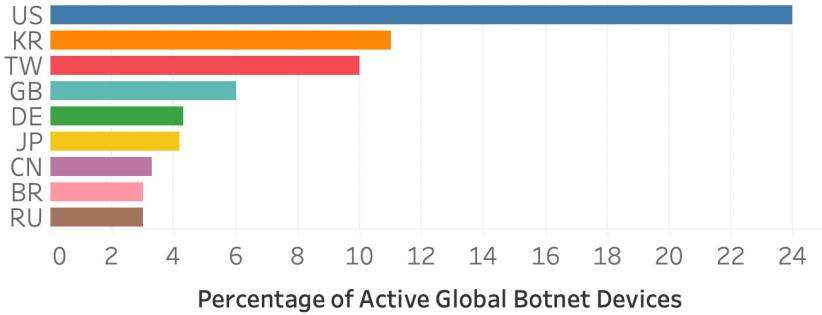

Source: Nokia Deepfield

How can we address this?

Traffic baselines! limiting!

How can we (really) address this?
#1 Anomaly detection

For >95% of DDoS, it's no longer about looking at what's inside the packet; instead, it's about who/what is sending the packet.

- bps/pps thresholds and baselines are insufficient and inadequate to track most of today's traffic (including flash crowd events)
- A **big data-driven approach** that correlates **network traffic** in real-time with **broader Internet context** (e.g., which type of device is behind a source IP address) is **much more effective** in reducing DDoS false-positive

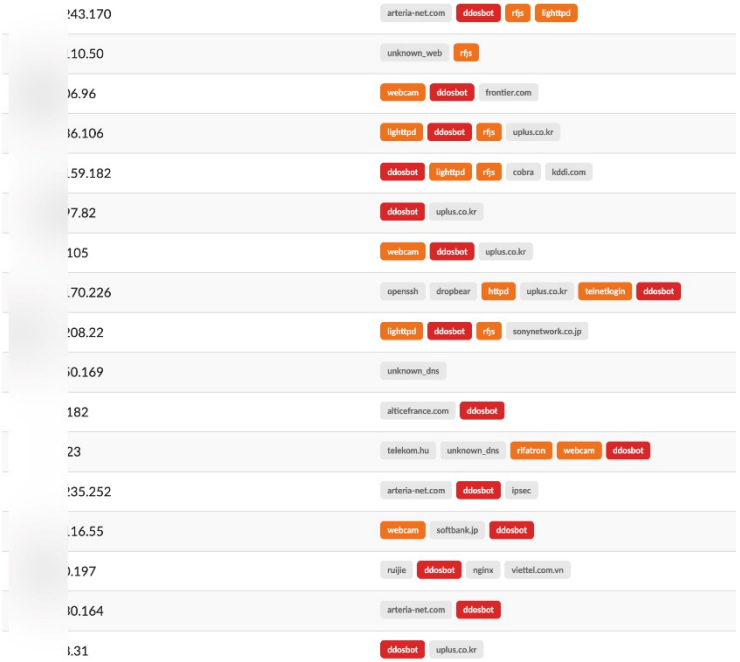

Nokia data: Top sources of traffic in DNS amplification attack towards a consumer IP address (target). Data source: GDTA-participating service and cloud providers globally, using the Nokia commercial DDoS defense solution.

How can we (really) address this?
Example Botnet against military network

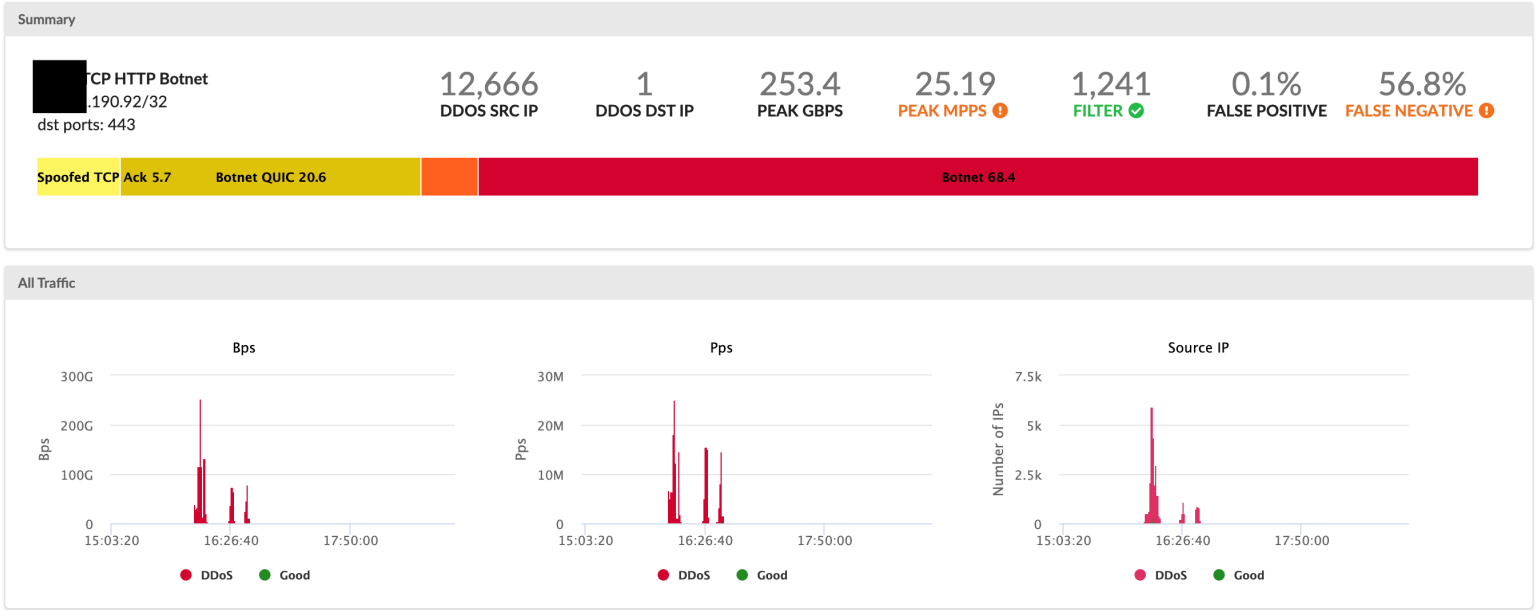

How can we (really) address this?
Example Botnet against military network

df["genome.src"].str.contains("edu") | Filter | Load all Flow | Show Detection | Show Genome Src

Time	TTL	Proto	TCPFlag	Peer	Src IP	SPort	Dst IP	DPort	Detect	Src Genome	Bytes	Len
16:03:40		17		[redacted]	[redacted]77.68	11965	[redacted]90.92	443	57 botnet_quic	webcam mini_httpd ddosbot [redacted].edu.ar	288598780	1,428
16:06:20		6	PA		[redacted]2.138	59182	[redacted]90.92	443	9 botnet	[redacted] uc-httpd ddosbot webcam	269323200	1,440
16:04:20		17			[redacted]77.68	11965	[redacted]90.92	443	57 botnet_quic	webcam mini_httpd ddosbot	160427230	1,428
16:36:20		6	PA		[redacted]77.68	54379	[redacted]90.92	443	9 botnet	webcam mini_httpd ddosbot	74304000	1,440
16:38:10		6	PA		[redacted]77.68	54379	[redacted]90.92	443	9 botnet	webcam mini_httpd ddosbot	67406400	1,440
18:41:10		6	PA		[redacted]0.84	54178	[redacted]90.156	5001		[redacted].edu	2304000	1,125
16:08:30		6	S		[redacted]23.175	36118	[redacted]90.219	80		nginx [redacted].edu	712680	60

How can we (really) address this?
xx.xx.77.68 is an NREN DDoS bot

Seen in multiple botnet DDoS attacks

mini_httpd 1.21 → vulnerable to CVE-2015-1548

How can we (really) address this?
#2 AI-based auto-mitigation

Once an attack is detected, a system can generate an automated response based on multiple parameters, which will create an optimized model for **that attack**, at **that time**, on **that network**.

For example:

- What's the attack vector mix?
- What mitigation devices are available on the network? At what scale and cost per bit?
- How can these devices be programmed?
- What's the botnet cluster launching that attack?

>95% of attacks can be mitigated on existing modern routers, thanks to progress on silicon performance, scale (e.g. 256k ACLs on Nokia FP4/FP5) and programmability (particularly NETCONF).

How can we (really) address this?
#2 AI-based auto-mitigation

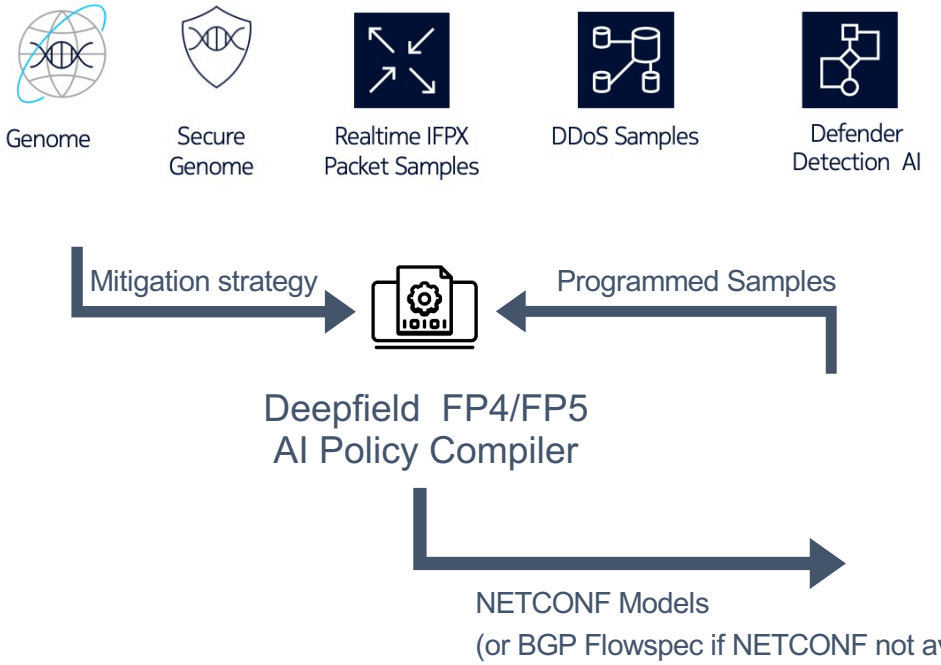


```
entry 8 create
    description ";#DFA;acl_90"
    match protocol 17
        dst-ip ip-prefix-list "VLAB_7_1"
        packet-length lt 40
        fragment false
    exit
    action
        drop
    exit
exit
entry 9 create
    description ";#DFA;acl_571"
    match protocol 6
        dst-ip ip-prefix-list "VLAB_7_1"
        tcp-fin true
        tcp-syn true
    exit
    action
        drop
    exit
exit
entry 10 create
    description ";#DFA;acl_579"
    match protocol 6
        src-ip ip-prefix-list "VLAB_9_518"
    exit
    action
        drop
    exit
exit
entry 4 create
    description ";#DFA;acl_13498"
    match
        dst-ip ip-prefix-list "VLAB_9_495"
        ttl range 1 37
    exit
    action
        drop
    exit
```

FP4

FP5

How can we (really) address this?
#3 Adaptive mitigation and collaborative learning

Instead of being driven by FUD:

- Mitigation effectiveness can be **measured** against the body of real-world attacks
- Model can be **trained** on new attacks to optimize countermeasures
- False-negative/false-positive rates can be understood and optimized

This requires **active collaboration between service providers**, to share (anonymized) DDoS threat intelligence data in near-real-time.

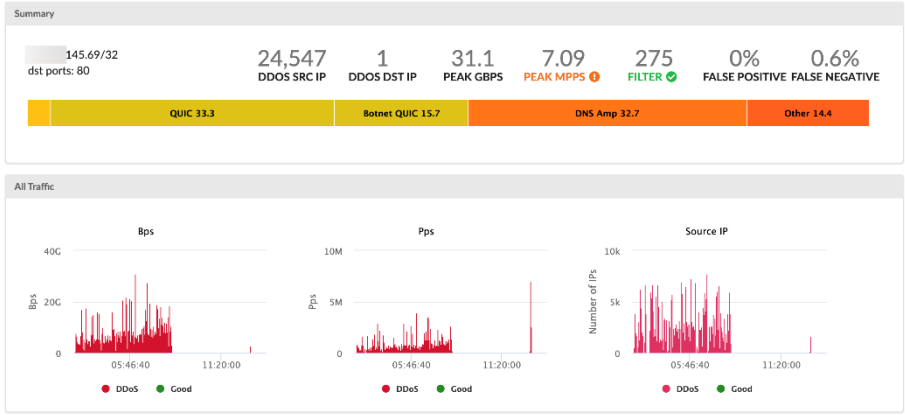

Summary from a DDoS attack in April 2023 on an EU government host.
Data source: GDTA-participating service and cloud providers globally, using the Nokia commercial DDoS defense solution.

Summary

DDoS botnets are nascent but already generate most of the DDoS traffic today

- Exponential growth of consumer & enterprise IoT
- ISP driving symmetrical 1Gbps connectivity further driving the “arms race”
- Nation-state attacks with large botnet networks

IoT botnets are everyone’s problem

- NRENs, ISPs, enterprises, vendors must take proactive IoT threat mitigation

AI/ML provides us tools to effectively address this threat

- Models can (and should) be trained on real-world data sets
- More collaboration is essential to share current DDoS data

tnc23
DIGITAL GENERATIONS
TIRANA, ALBANIA | 5-9 JUNE 2023

Thank you

Any questions?

GÉANT
Co-funded by the European Union