



Building a future-proof Cyber Fusion Center

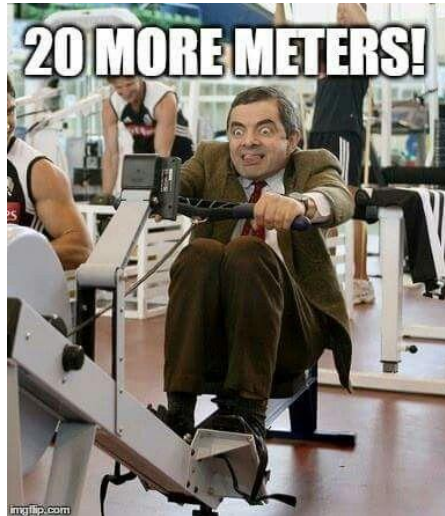
CIO/CISO Future Connections München

Maxim Deweerdt
NVISO



Introduction

Who am I and what do I want to talk about?



Finding the right SOC partner

Balluff's journey – Customer Reference

BALLUFF

Emanuel Somosan

Manager Global Security Operations

Introduction

Who am I and what do I want to talk about?

I want to talk about how the **cyber fusion center of the future** is built.
Before that, we will introduce some of the common problems of traditional SOC's:



Expensive



Alert Fatigue



Skill Shortage



Ever-Expanding
Landscape

Critical Success Factors

How to build a highly functioning Fusion Center



Threat-Centric



Purple Focus



SOAR-Centric

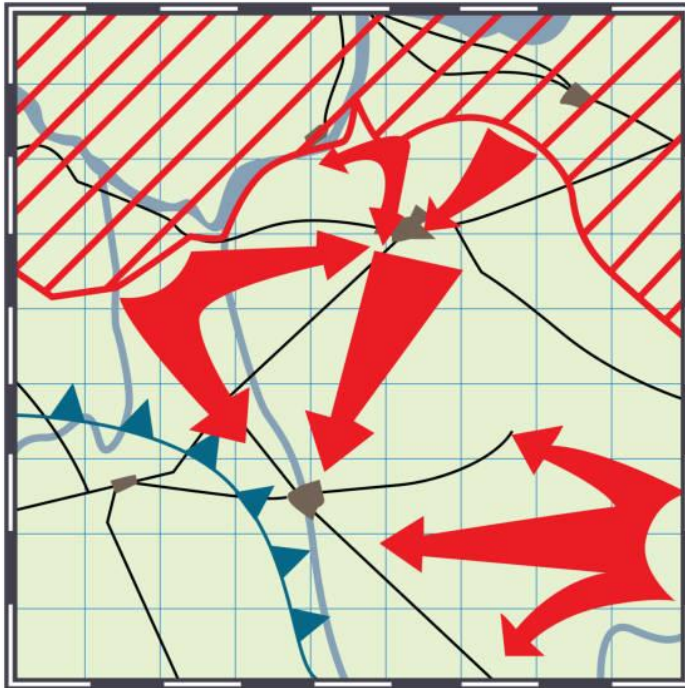


Automation-First



Threat-Centric

How do we win? Adversaries have goals in mind





Threat-Centric

Leveraging MITRE ATT&CK for a threat-centric approach

ATT&CK for Adversary Emulation

When organizing adversary emulation (such as Red or Purple Team exercises), the emulation plan can be based on MITRE ATT&CK. This facilitates tracking & reporting.

ATT&CK for Threat Intelligence

When consuming or generating Threat Intelligence, observed adversary behavior can be mapped to MITRE ATT&CK. Several platforms support this mapping (e.g., MISP has a MITRE ATT&CK mapping).

MITRE ATT&CK should be the **common language** for the Cyber Fusion Center

ATT&CK for Detection Capability

The overall detection capability of an organization can be mapped to MITRE ATT&CK. This facilitates, for example, reporting on the maturity / scope of the SOC.

ATT&CK for Defense Prioritization

In addition to measuring the detection coverage using MITRE ATT&CK, we can do the same for preventive controls. What MITRE ATT&CK techniques do we actively block?



Purple Focus

Combining Red and Blue skills

EXAMPLE OF BAD “RED” THOUGHTS

Report with many vulnerabilities = well done!

Success is measured by # of failed controls

No big incentive to help Blue Team, as Blue Team failure = Red Team success!

EXAMPLE OF BAD “BLUE” THOUGHTS

No alerts = Wow, our preventive controls are working really well! 😊

A lot of alerts = Wow, we have a good coverage of detections

No big incentive to help Red Team, as Red Team failure = Blue Team success!



Purple Focus

Combining Red and Blue skills

The Cyber Fusion Center should be **a purple ambassador** and make sure red thinks a bit more blue, while blue should think a bit more red:



**Red Team with a
“touch of blue”**

- **Understand prevention, detection, and response techniques**
- **Understand complexities** and limitations of target organization and tailor recommendations
- **Present known TTPs** to Blue Team (highlight "quick wins") and innovate Red Team approach continuously



**Blue Team with a
“touch of red”**

- Understand and follow up on known adversary TTPs
- **Test individual TTPs continuously** and improve where possible
- Track and report on **coverage of TTPs** (e.g., ATT&CK framework)



Purple Focus

Combining Red and Blue skills

So... No more yearly red teams? There's room for both:



Red Team

Organize **periodic Red Team exercises** to **assess** the actual state of security in the organization. Offer feedback only after the exercise ends, as the exercise is typically meant to be stealthy (realistic adversary emulation)...

VALUE: Periodic assessment of organization resilience



Purple Team

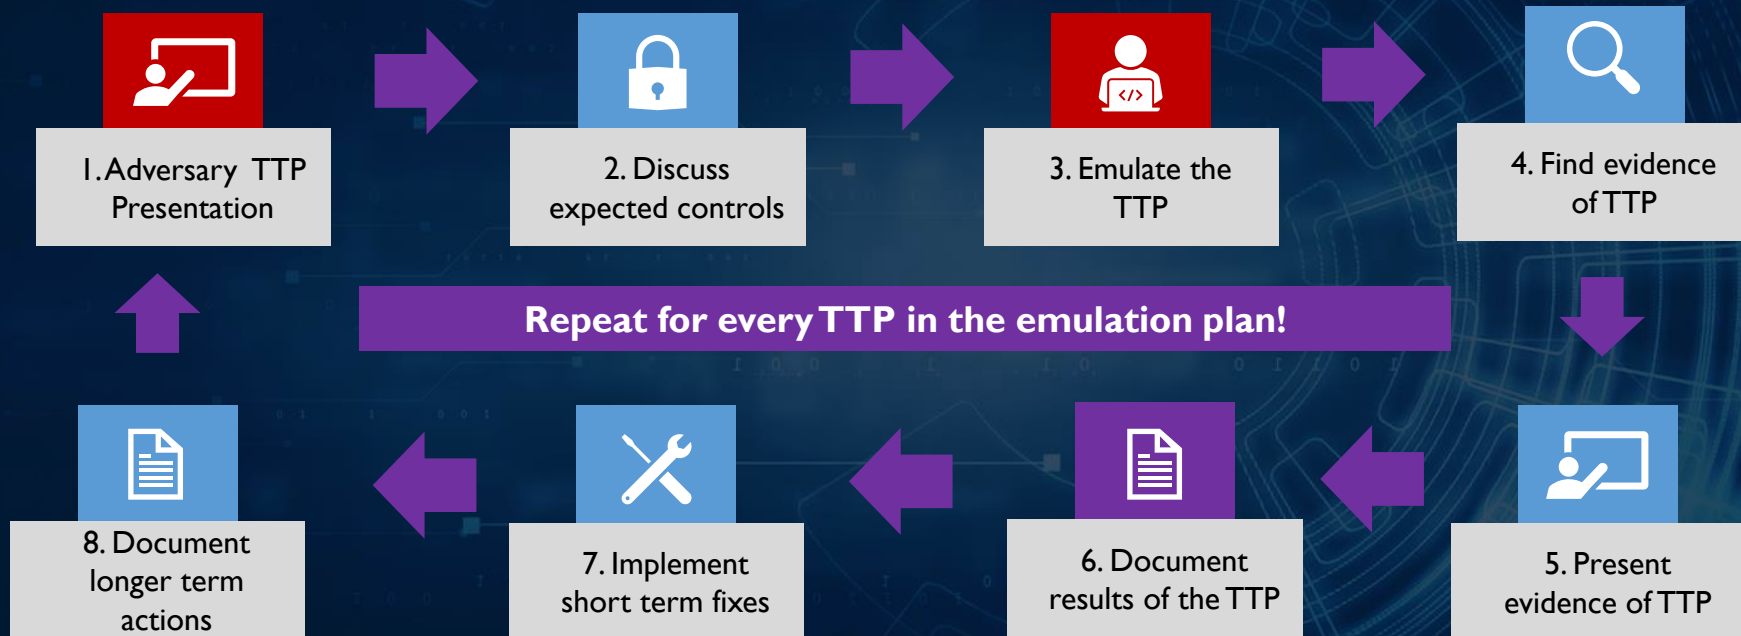
Perform **continuous Purple Teaming** to improve the state of security in the organization. Blue Team members simulate focused attack techniques as part of their operations to immediately test effectiveness of detection and prevention controls.

VALUE: Continuous improvement of organization resilience



Purple Focus

Organising a Purple Team Exercise

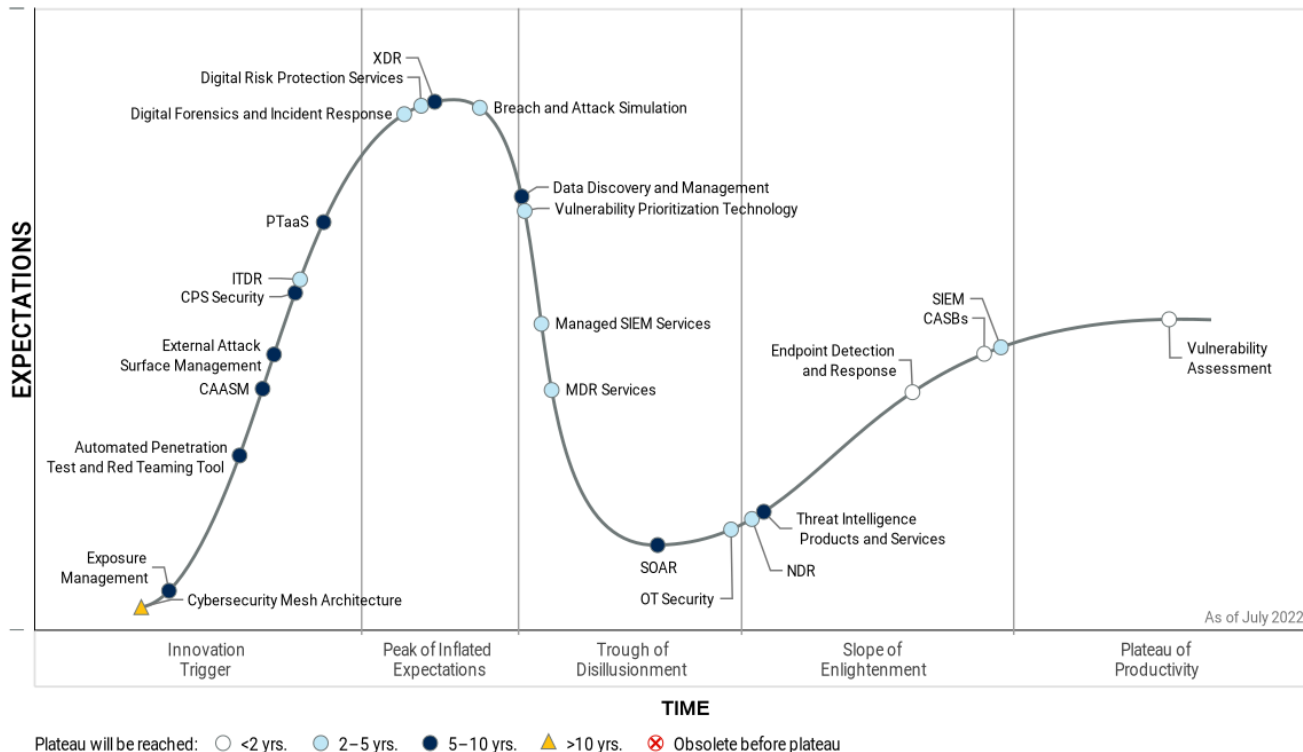




Purple Focus

Continuous validation of security controls / detection rules

Hype Cycle for Security Operations, 2022



Breach and Attack Simulation (BAS) tools are designed to support organizations in **automatically emulating techniques and attacks** leveraged by real threat actors.

They are **not designed to replace a red team** (as they lack the manual evasion techniques the red teamers will apply), but can be tremendously helpful to **continuously test defenses against a predefined attack chain**.



SOAR-Centric

Security Orchestration, Automation & Response

SOAR

Security Orchestration, Automation and Response (SOAR) tools refer to a collection of tools that help organizations coordinate, execute and automate tasks between security tools and people. They are composed of 4 main blocks:



Integration with security tools with plugins to build **security playbooks** to automate tasks and respond to alerts automatically.



Present **contextualized** information and **enriched** alerts to allow analysts to **take decisions and actions quickly**.



Provide **reports and insights** about manual and automatic actions and about possible improvements.

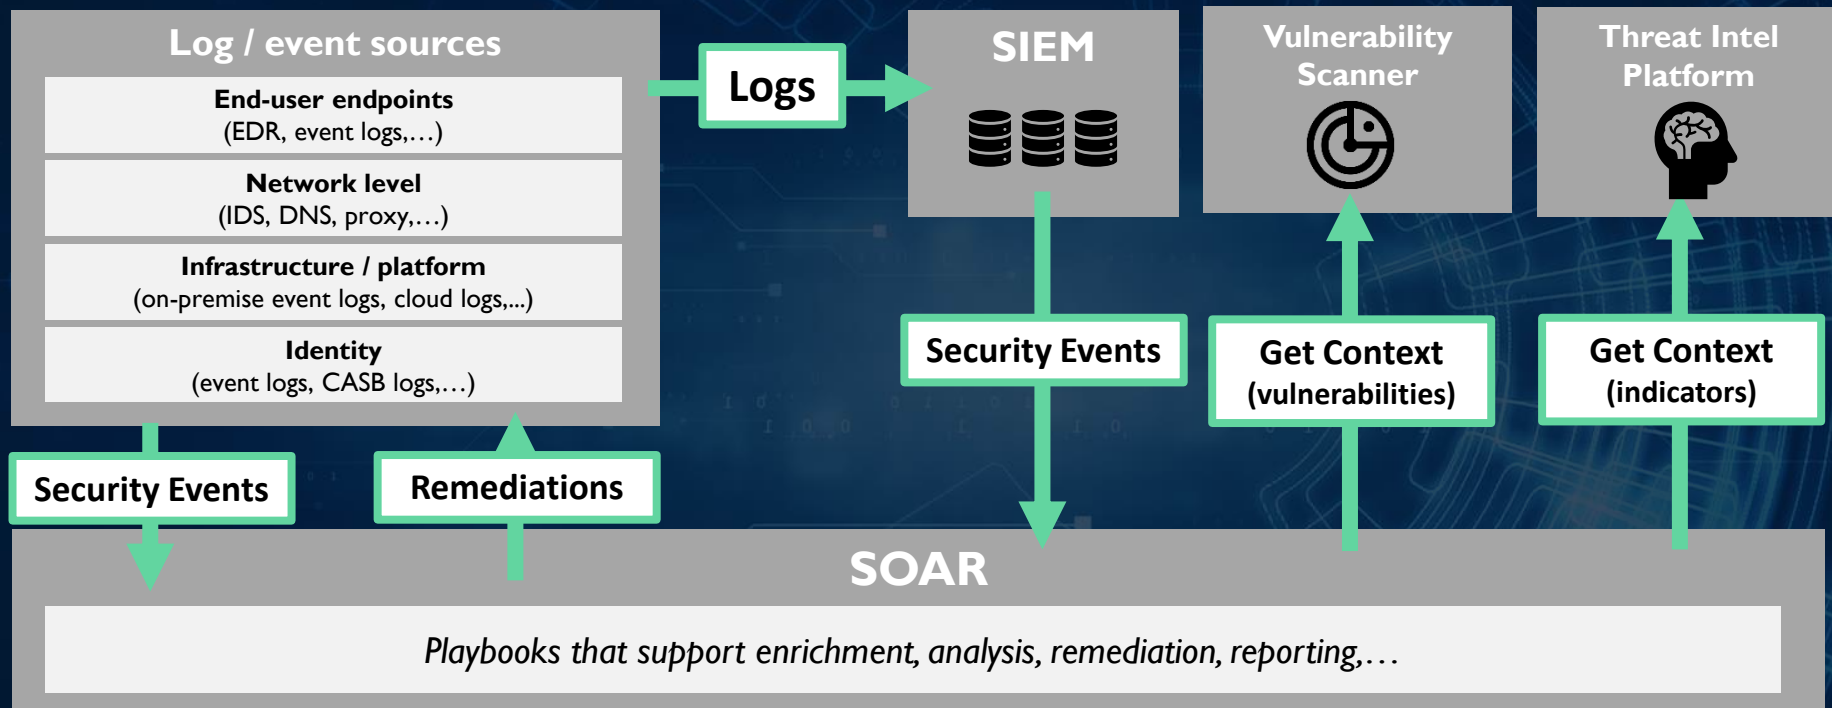


Provide one **single centralized platform** for analysts with all the **dashboards** and **alerts** from the different sources.



SOAR-Centric

A SOAR-Centric architecture



The **SOAR platform becomes the “central brain”** of the Fusion Center (instead of the SIEM). All security technologies should be connected to the SOAR (both for detection, contextualisation, handling, reporting and remediation)



SOAR-Centric

Marriage between automation and human effort

- 1 A user reports a suspected phishing email (e.g. using Outlook button)
- 2 Automatic check for URLs and Indicators of Compromise (IoCs)
- 3 Based on available data / context, make decision on benign / malicious
- 4 If confirmed malicious, scan the user endpoint for malware
- 5 Block incoming e-mails with similar properties (URLs, sender, subject,...)
- 6 Automatically remove already delivered e-mails from mailboxes
- 7 Provide feedback to reporter + warn others about the phishing attack





Automation

Automation is a vital component

L1 Security Analyst Industry Stats:



20 Minutes Per Security Event



25 Events Per Day

Industry lacking two million personnel worldwide so how does an operation cope?

Automation is a key component



647 Security Events = 26 Analysts

24x7 = 12 Analysts Minimum

Automation has decreased the analytical workload by **97%**.

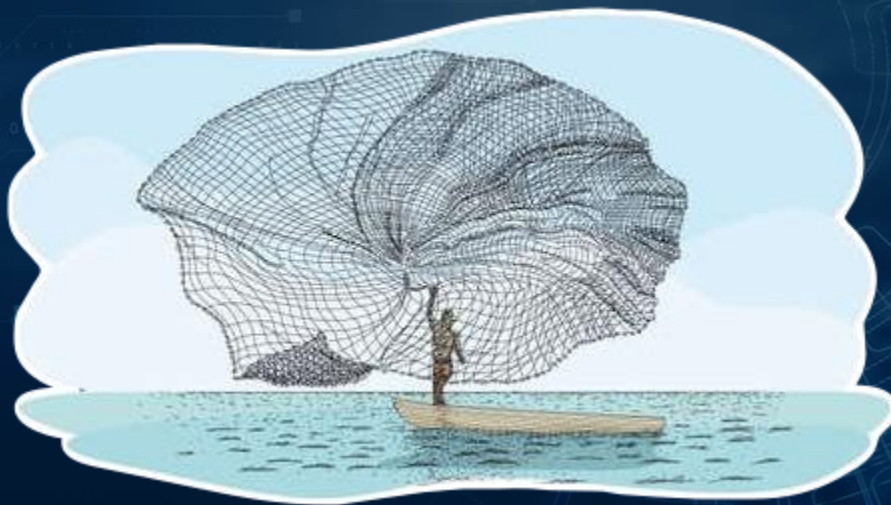
Of **5.790** alerts only **145** were manually analysed in 7 days in the SOC. This reduced costs down to 1.500 EUR as opposed to **60.000** EUR.



Automation

So how can we do this right?

Cast a (relatively) **wide net**, but **optimise (automate) incident analysis** and response!





Automation

So how can we do this right?

“Geographically improbable log-on for user Maxim Deweerdt”



Automation

So how can we do this right?

“Geographically improbable log-on for user Maxim Deweerdt”

Enrich: Add privileges of user Maxim Deweerdt to security event

Enrich: Add insights & reputation of source IP address to security event

Enrich: Add whether or not MFA was used in authentication to security event

Enrich: Add historic locations used by Maxim Deweerdt to security event

Enrich: Add security risk score for user Maxim Deweerdt to security event

Enrich: Add info on workstation security alerts for Maxim Deweerdt’s workstation to security event

Enrich: ...

Decide: Confirm whether, based on the above enrichments, a false positive can be confirmed

Remediate: When confirmed true positive (and allow-listed for remediation), execute remediation

Present: When unsure, present enriched security event to analyst for further follow-up





Automation

So how can we do this right?

Indicators (4)

Indicators

Type	Value	Verdict	Related Incidents	First Seen	Last Seen	
IP	20.223.215.19	Benign	2	February 8, 2023 17:43	11@340439	
User Agent	<div><h3>Closing Information</h3><div><div>Closed Time</div><div>February 8, 2023 17:44</div></div><div><div>Extended Close Reason</div><div>False Positive</div></div><div><div>Close Notes</div><div>guillaume@qa-nviso.be generated an access anomaly from the IP: 20.223.215.19</div><div><ul style="list-style-type: none">• All of the incident-involved IPs leveraged MFA at least once for the logins attributed to this incident, resulting in it being considered a false positive.</div><div>These findings justify our assessment of this Incident being a False Positive</div></div></div>					11@340439
Account						11@340439
Domain						211@380488
Azure AD Login Overview						
IP Address	Country	City	OS	Browser	Count	
109.134.85.71	BE	Brussels	Windows 10	Edge 109.0.1518	1	
109.134.85.71	BE	Brussels	Windows 10	IE 11.0	0	
84.199.226.10	BE	Brussels	Windows 10	Edge 109.0.1518	0	
84.199.226.10	BE	Brussels	Windows 10	Edge 109.0.1518	0	
146.103.254.16	BE	Leuven	Windows 10	Edge 108.0.1462	0	
109.134.85.71	BE	Liege	Windows 10	IE 11.0	1	
20.216.172.231	FR	Paris	Windows 10	Edge 108.0.1462	0	
20.223.215.19	IE	Dublin	Windows 10	Edge 109.0.1518	0	



Automation

So how can we do this right?

Investigation Data

Details

REDACTED, and REDACTED generated access anomalies from the IPs: 162.241.87.65,

192.168.59.86, 62.134.91.130

• Some of the observed IPs

the involved user: DE: 62.134.91.130

Based on these findings, w

seen before for

Automated Remediation

Revoke sessions, tokens and
reset credentials

Azure AD Login Overview

IP Address	Country
62.134.91.130	DE
62.134.91.130	DE
62.134.91.130	DE
62.134.91.130	DE
162.241.87.65	US

Neuhausen Auf Den Fildern

Windows 10

Edge 18.19044

25

16

5

4

Neuhausen Auf Den Fildern

Windows 10

Edge 109.0.1518

21

21

0

0

Provo

Windows 10

Chrome 69.0.3497

5

0

5

0

MFA OK	MFA Fail	SFA OK
94	3	0
0	0	55



Automation

So how can we do this right?

Last 30 days

TP/FP

TIME SAVED WITH
AUTOMATION

3.500



Access Anomaly events

200 True Positive

3.300 False Positive



Automatically
analyzed and
closed


66.000 minutes / 1.100 hours /
46 days

(Average of ~20min per analysed
event)

AI as a force multiplier

Leveraging AI as an internal knowledge base




Chat Readme + New Chat  


 **Chatbot**

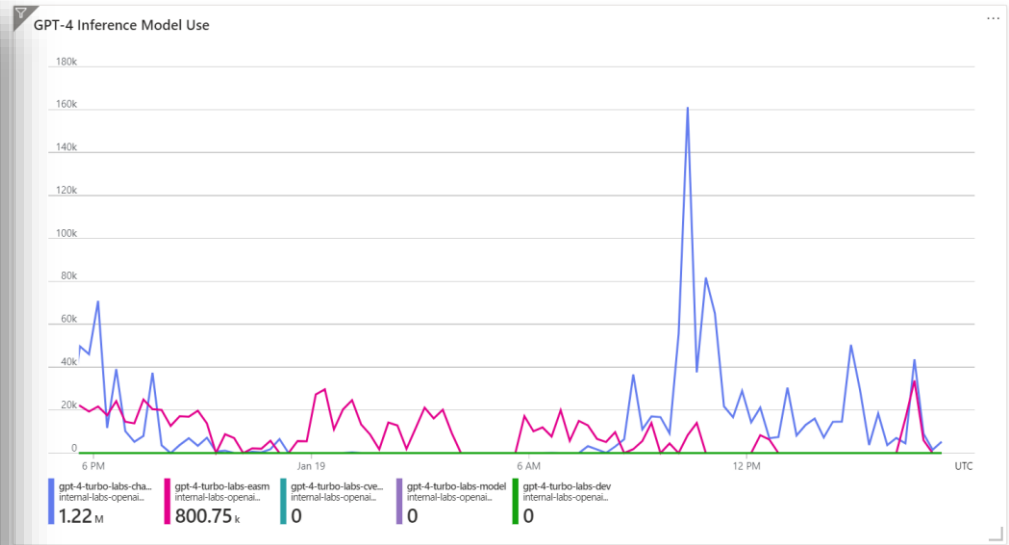
All conversations are handled within our own EU Azure environment, and are not shared with OpenAI.
⚠️ Please respect the [Code of conduct for Azure OpenAI Service](#) and don't get us banned ❤️

Use the '+' button below to upload any file and ask questions about it.
Write `RFP` command below to generate answers to respond to Security Questionnaires.
Write `GP` command below to forget about NVISO & answer with general-purpose knowledge only.

[Product Roadmap](#) | Questions or ideas on how to improve this Chatbot? [Submit a ticket!](#)
GPT-3.5 Turbo running in Azure EU, trained with NVISO-specific information, as well as internet information up to September 2021.

  Type your message here... 

Built with  Chainlit



How to let your organization use GPT without the privacy/security/ethics risk? **Deploy your own GPT!**

AI as a force multiplier

Leveraging AI to analyze phishing emails

Field	Type
_expires	datetime
_query_time	float
_reputation	str
_updated_at	datetime
category	str
confidence	int
explanation	str
prompt	str
risk_score	int
source	str

You are a language model helping a security analyst to decide if an email is phishing, or not. What follows is the full email that was sent to our phishing analysis mailbox by the recipient of the email.

Your response should be a valid JSON structure containing the following fields:

is_phishing: Yes or No

confidence: High, Medium or Low. This indicates how confident you are in your decision.

[...]

One of the key examples of how we use OpenAI in the SOC is the analysis of phishing e-mails.

How you “prompt” the AI is however of the utmost importance!

AI as a force multiplier

Leveraging AI to enhance incidents/tickets

Insights from our NITRO MDR Copilot (Beta)

Description of the incident:

On September 26th at 17h01 UTC, DNS transactions from your resource were analyzed and compared against known malicious domains identified by threat intelligence feeds. As a result, communication with [redacted] was compromised.

Potential Causes:

There are several possible causes for this incident. Another possibility is a misconfiguration in your network that is causing the communication to be blocked.

Potential Risks:

The potential security risk is that the compromised communication could lead to a data breach or a loss of sensitive information.

Proposed Mitigation:

1. Investigate the communication logs for any suspicious activity.
2. Verify that your network configuration is correct.
3. Implement network security measures to prevent future incidents.

You're an AI assistant for the SOC that will help by enriching Security incidents tickets. Those tickets are automatically generated in XSOAR and sent to Jira. Those tickets are not very readable for humans and contain technical information. Your job is to enrich the tickets with additional information explaining the potential causes of the alert, the potential risks involved and a few potential steps to mitigate.

Use the following structure:

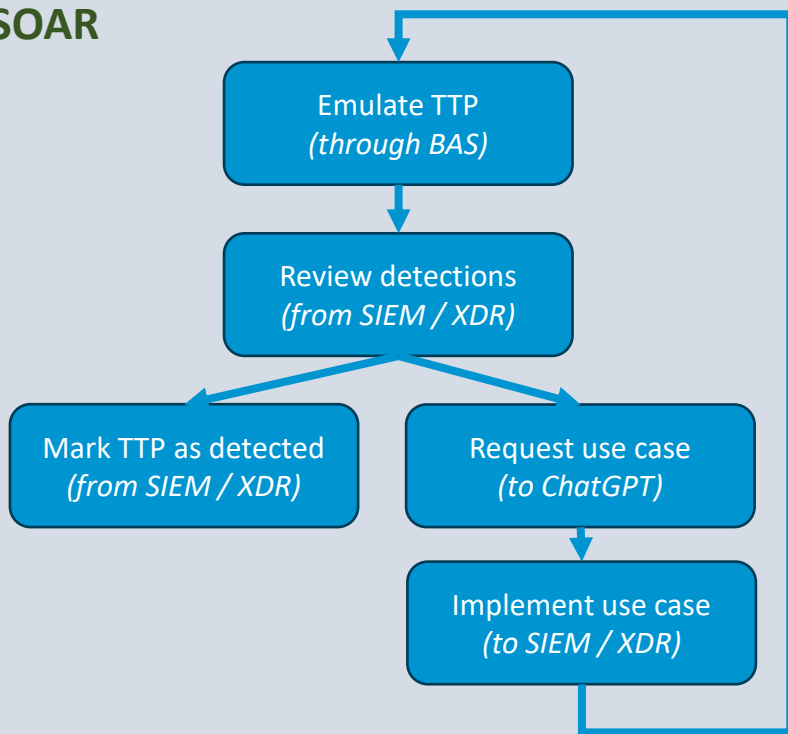
[...]

Add ChatGPT-generated insights to tickets when communicated to your constituents.

AI as a force multiplier

Leveraging AI in the detection engineering pipeline

SOAR



An active topic of research at Nviso is how we can further automate **detection engineering**. High-level action plan is below:

1. Trigger emulation of a TTP by calling the BAS (Breach Attack Simulation) tool
2. Review detections from SIEM / XDR (which are already being ingested in the SOAR platform) and determine whether the TTP was successfully detected
3. If the TTP was successfully detected, mark it as such and move to the next TTP.
4. If the TTP was not successfully detected, generate a detection analytic specific to our technology through ChatGPT
5. Push the suggested detection analytic to the SIEM / XDR for testing
6. Rinse and repeat

AI as a force multiplier

Leveraging AI in the detection engineering pipeline

You
Create an EQL query that detects the MITRE technique T1087. Return only the query.

ChatGPT

arduino

Copy code

```
process where process_name in ("net.exe", "net1.exe", "nltest.exe") and command_line
```

You
Create an EQL query that detects credential dumping. Return only the query. Show it as text not code.

ChatGPT

```
process where process_name in ("mimikatz.exe", "lsass.exe", "procdump.exe", "dumpert.exe") and  
(parent_process_name in ("svchost.exe", "services.exe", "explorer.exe", "lsass.exe", "winlogon.exe"))
```

Conclusion

How to build a highly functioning Fusion Center



Threat-Centric

GOAL ORIENTED DEFENSE

Only way to achieve upper hand and make the blue team “win”.



Purple Focus

CONTINUOUS TESTING

Continuous improvement of visibility gap and of post-compromise detection opportunities.



SOAR-Centric

HEART OF YOUR OPERATION

Stop leveraging the SIEM as the heart of your SOC – it was not made for this.



Automation-First

STOP “MONKEY” TASKS

Only way to retain skilled people and scale your operation.

Building the Cyber Fusion Center of the future

Q&A

Q&A



Maxim Deweerdt
maxim.deweerdt@nviso.eu