



# Microsoft Security Copilot – ako pomôže bezpečnostým tímom

Roman Tuchyňa | ATS| MS Slovakia



# Benefits of AI for security

- > **Efficiency:** Prioritization and automation
- > **Speed:** Ability to understand unique threats in real time
- > **Scale:** Ability to process large volumes of data

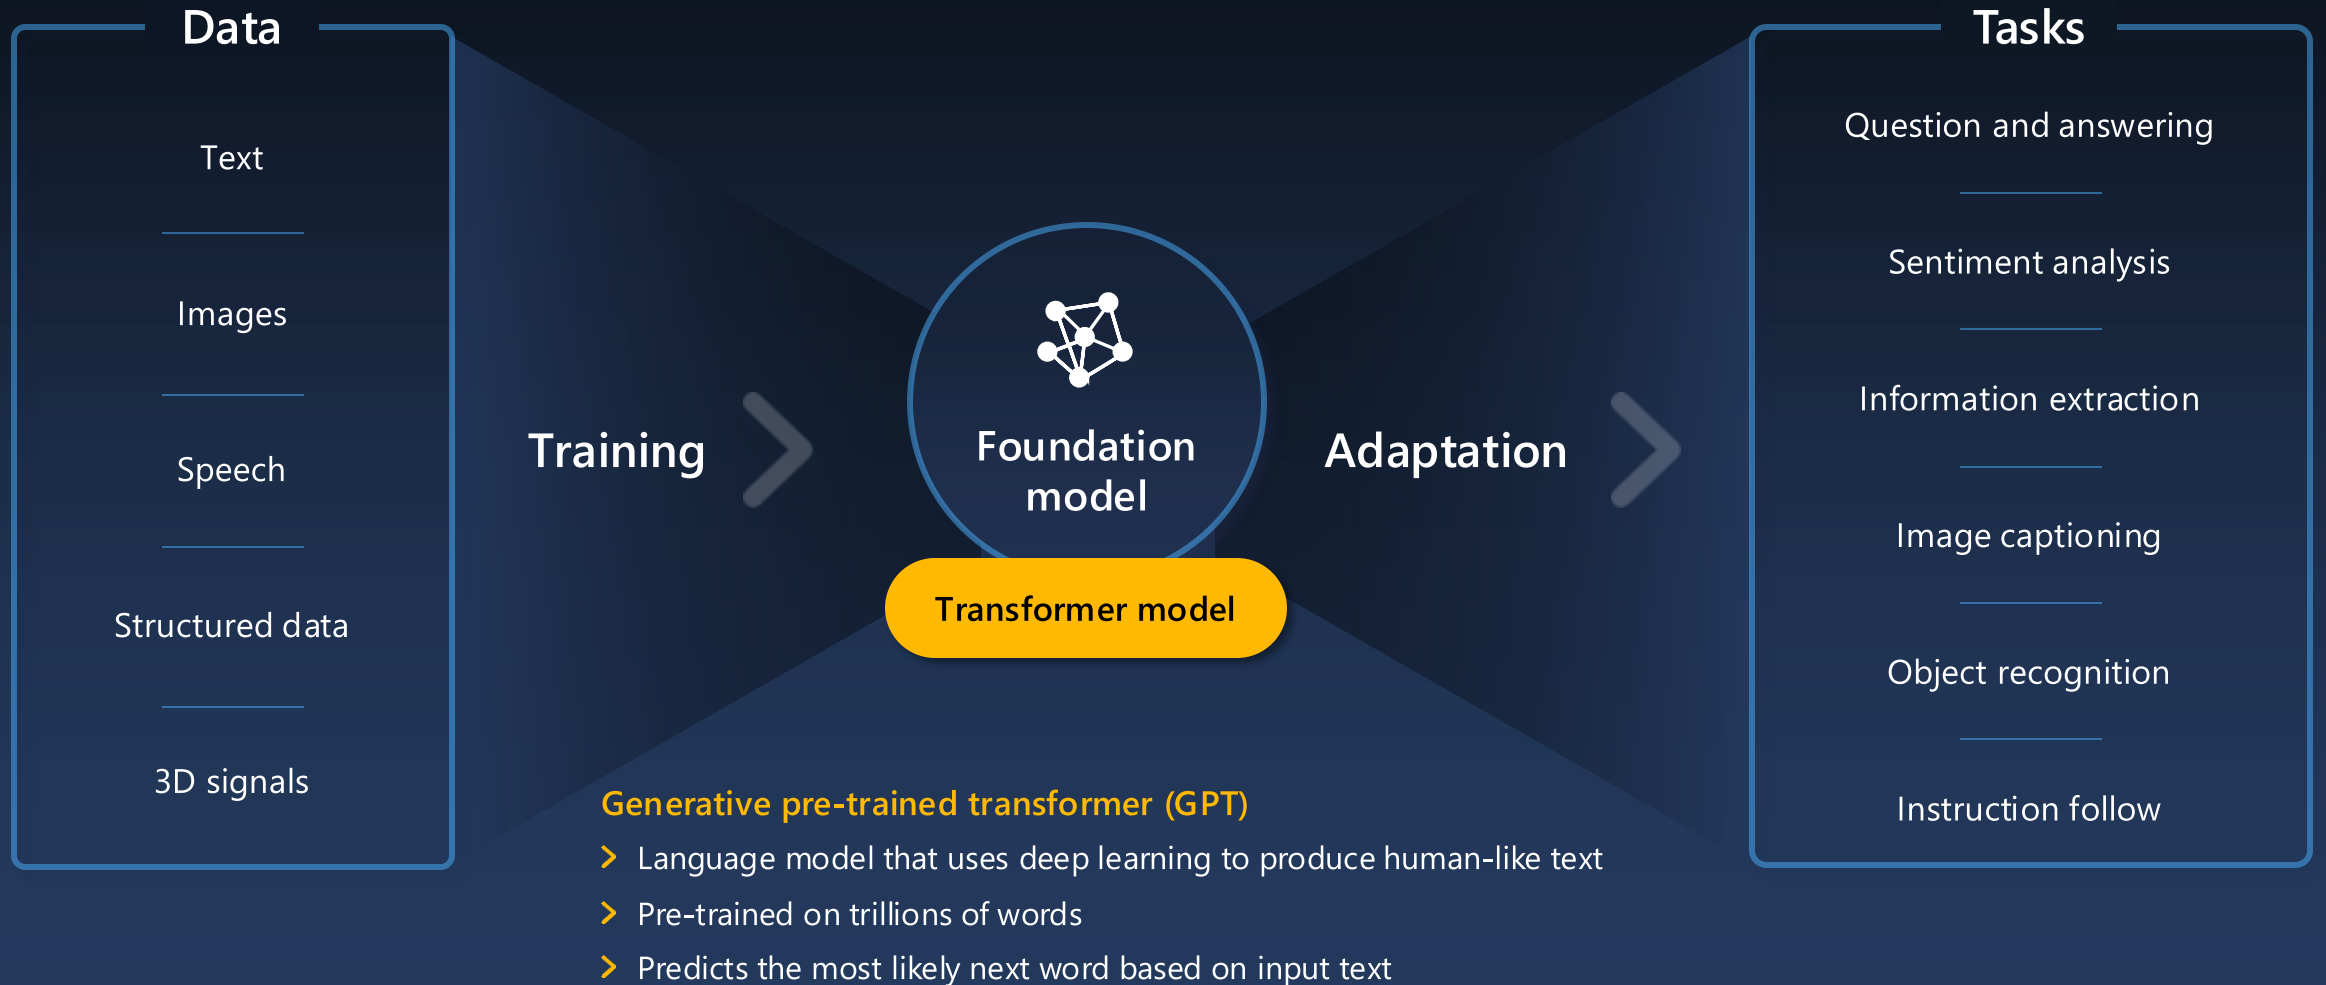


The Microsoft Sentinel platform has more than **10 petabytes** of daily ingestion

What makes  
Generative AI important for  
security?



# Understanding foundation models





# Modernizing your security operations



## Security posture management

Discover whether your organization is susceptible to known vulnerabilities and exploits. Prioritize risks and address vulnerabilities with guided recommendations.



## Incident response

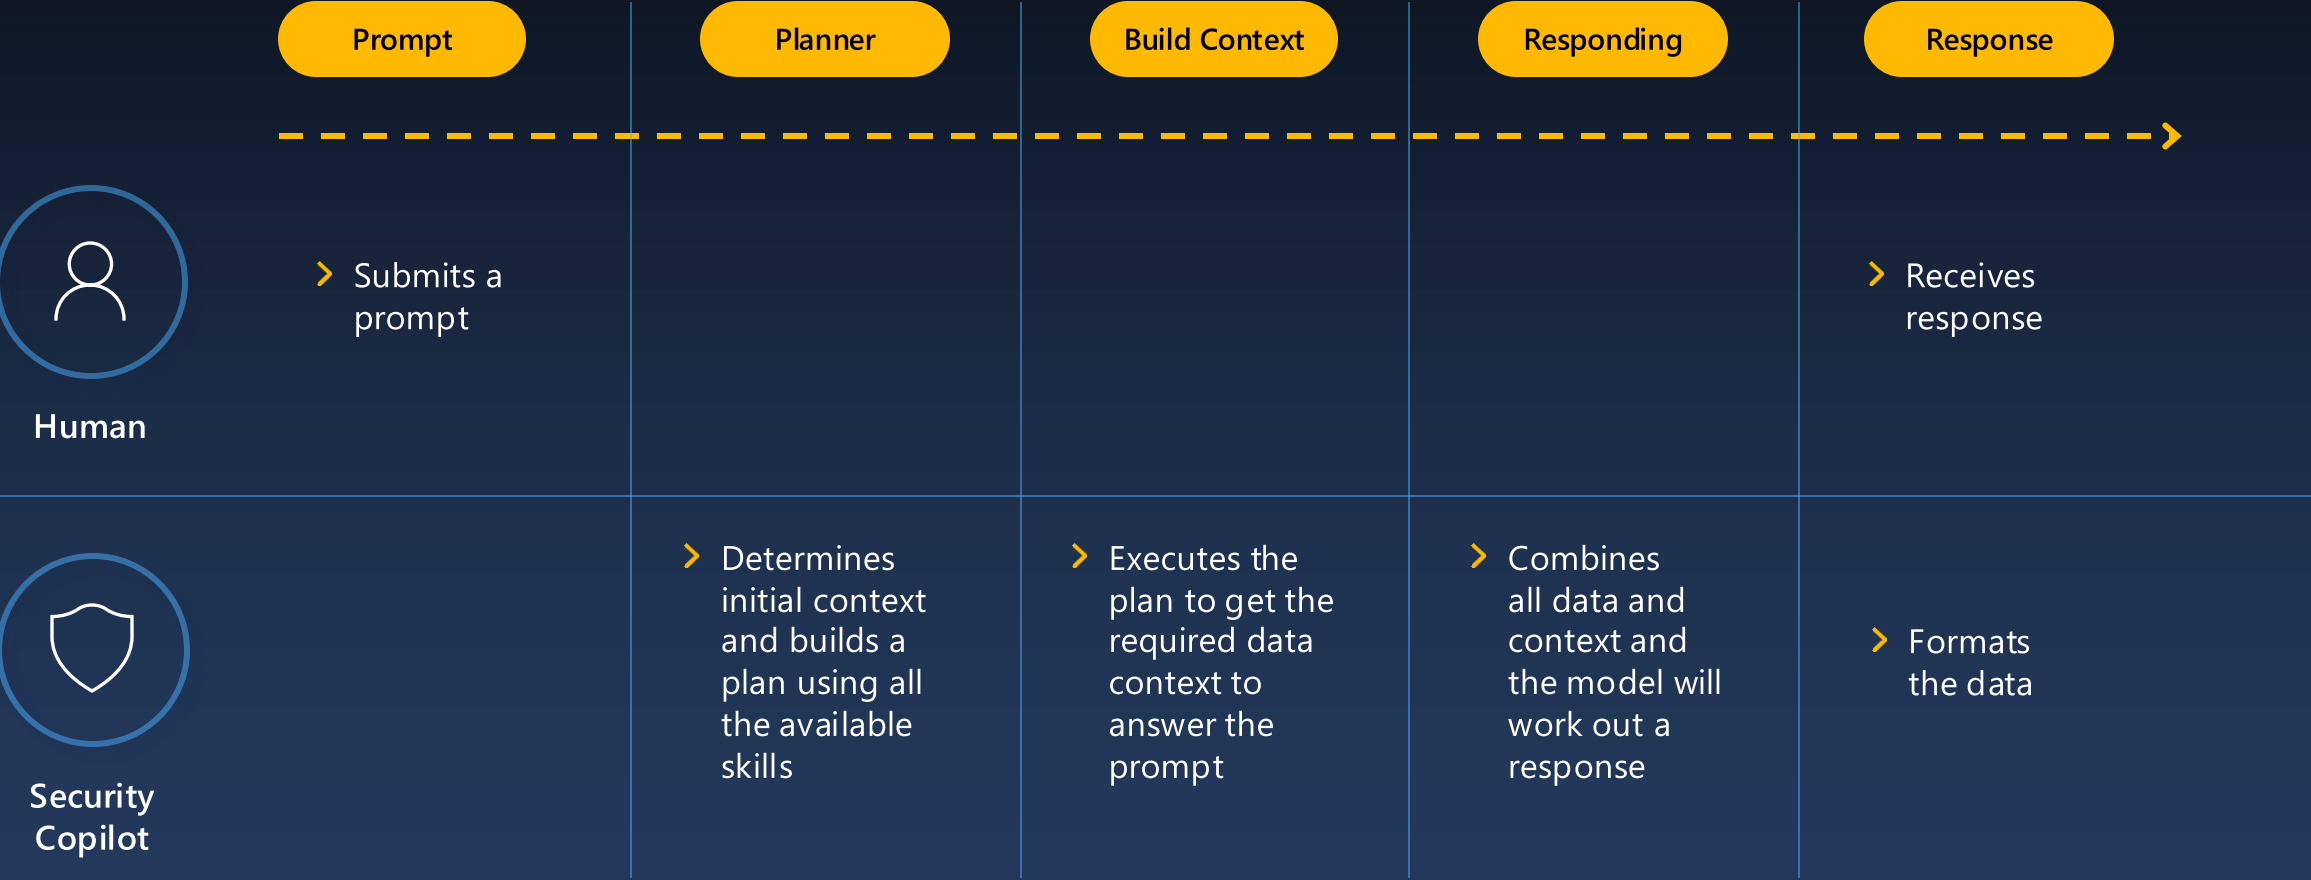
Surface an ongoing incident, assess its scale, and get instructions to begin remediation based on proven tactics from real-world security incidents



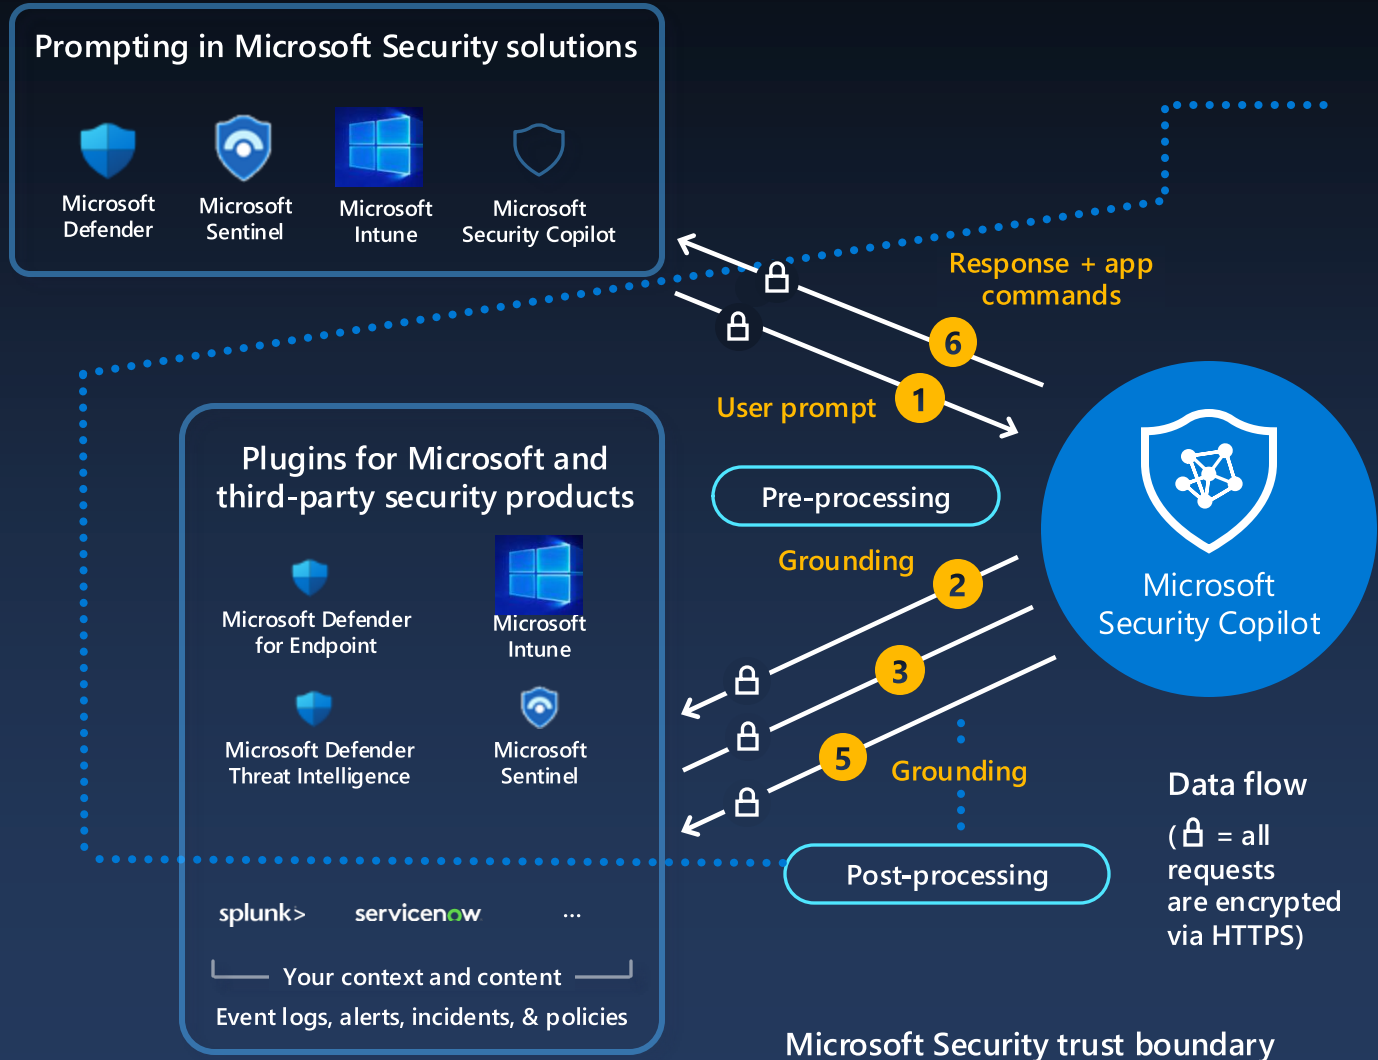
## Security reporting

Summarize any event, incident, or threat in seconds and prepare the information in a ready-to-share, customizable report for your desired audience

# How it works



# Data flow for Microsoft Security Copilot



💠 Tell me about my latest incidents.

💠 Are any of my machines affected?

💠 Summarize this incident in PowerPoint.

💠 Which alerts are being triggered the most?

💠 What is log4shell?

💠 How can I improve my security posture?

How can I improve my security posture

Jan 10, 12:38 PM

Here are three evidence-based recommendations to improve your security posture:

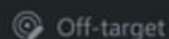
1. **Use Microsoft Intune ASR rules to prevent common attack techniques:** This will help protect your devices from common attack techniques. Currently, 438,174 out of 758,659 devices are impacted. Learn more about ASR rules [here](#).
2. **Use Microsoft Intune to enforce the latest Defender security baseline:** This will ensure that your devices are up-to-date with the latest security settings. Currently, 438,174 out of 758,659 devices are impacted. Learn more about the Defender security baseline [here](#).
3. **Use Microsoft Intune Endpoint Privilege Management to reduce exposure by enabling standard users and limiting local admin accounts:** This will help minimize the attack surface by reducing the number of local admin accounts. Currently, 438,174 out of 758,659 devices are impacted. Learn more about Endpoint Privilege Management [here](#).

▼ Sources

Microsoft Intune



Confirm



Off-target



Report



Pin



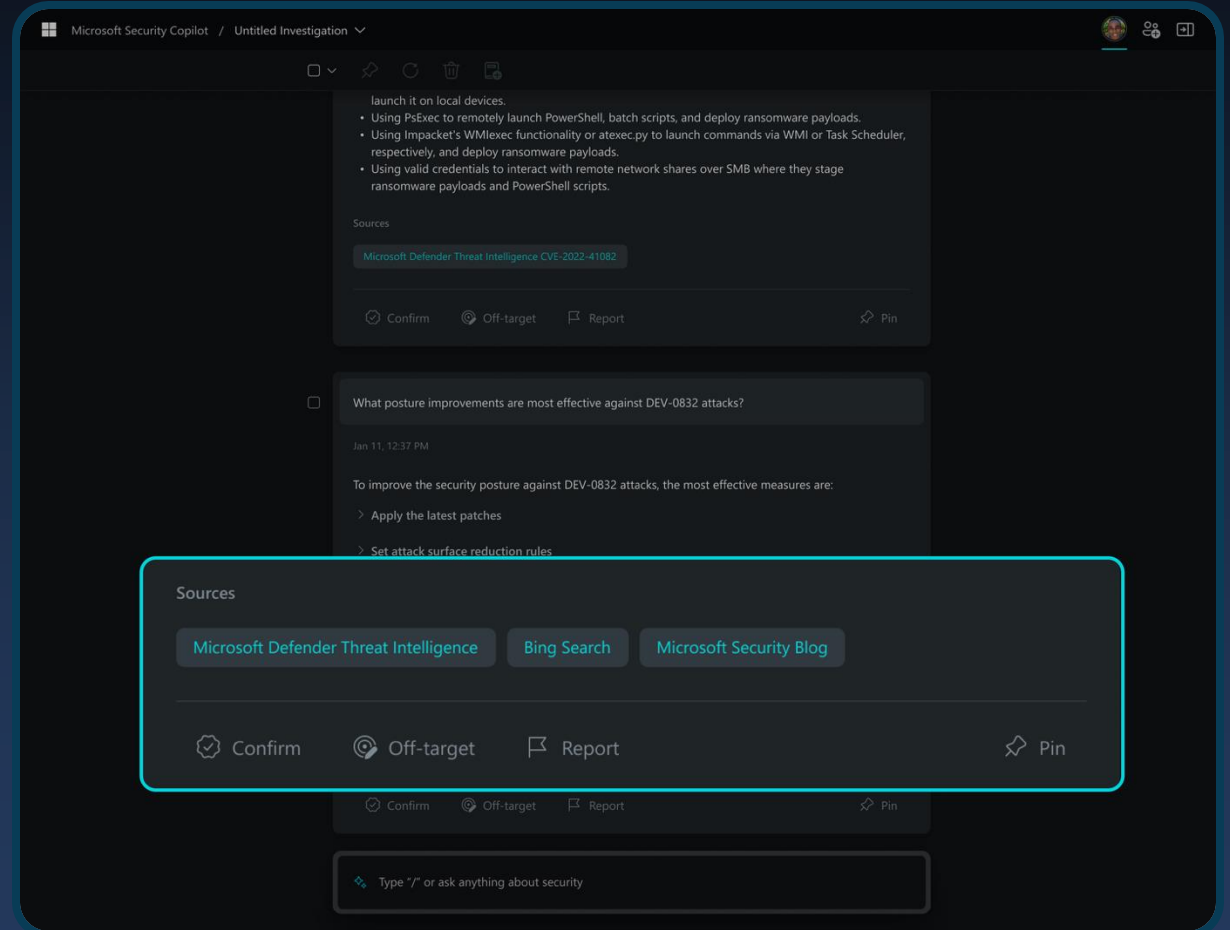
# Building trust with AI

Trust is fundamental to a healthy relationship and if Copilot is meant to augment a human, then we must find ways to build Trust between man and machine.

One hinderance is hallucinations. A hallucination is generated content that appears plausible but is either factually incorrect or unrelated to the provided context. It comes across as qualified knowledge, wrapped in a confident response - aka =Bull\*\*it

## Impact

1. Show reasoning, sources, debug and runtime
2. Ensure data is compliant, secure, and private
3. Address harms and hallucinations
4. Be transparent and allow for an open dialog



# Prompting is not chat

We are leveraging prompt-based experiences that differ from back-and-forth “chat” conversations. We consider prompts to be natural language programs interacting with the model to get accurate results that help optimize and define workflows.

The impact of freeing ourselves from existing thinking pushed us in new directions.

## **Impact**

1. New paradigm that feels familiar
2. Less question and answer, more like a coworker who does the work
3. Notebook style context-based interaction model
4. Investigation as a natural language notebook

# Put the user in control

AI is built on probabilities and will make mistakes, so we need to design for it being wrong. Find ways to always keep the human in control. Allow the human to decide what is important, what is relevant and what isn't. Focus on the human to be the one that takes action.

This will reduce an overreliance on AI and will build trust and confidence.

## **Impact**

1. Allow users to control and grade the AI output
2. Give a user tools to edit and correct AI outputs
3. Build affordances for providing feedback

# Repeatable tasks are now bundled

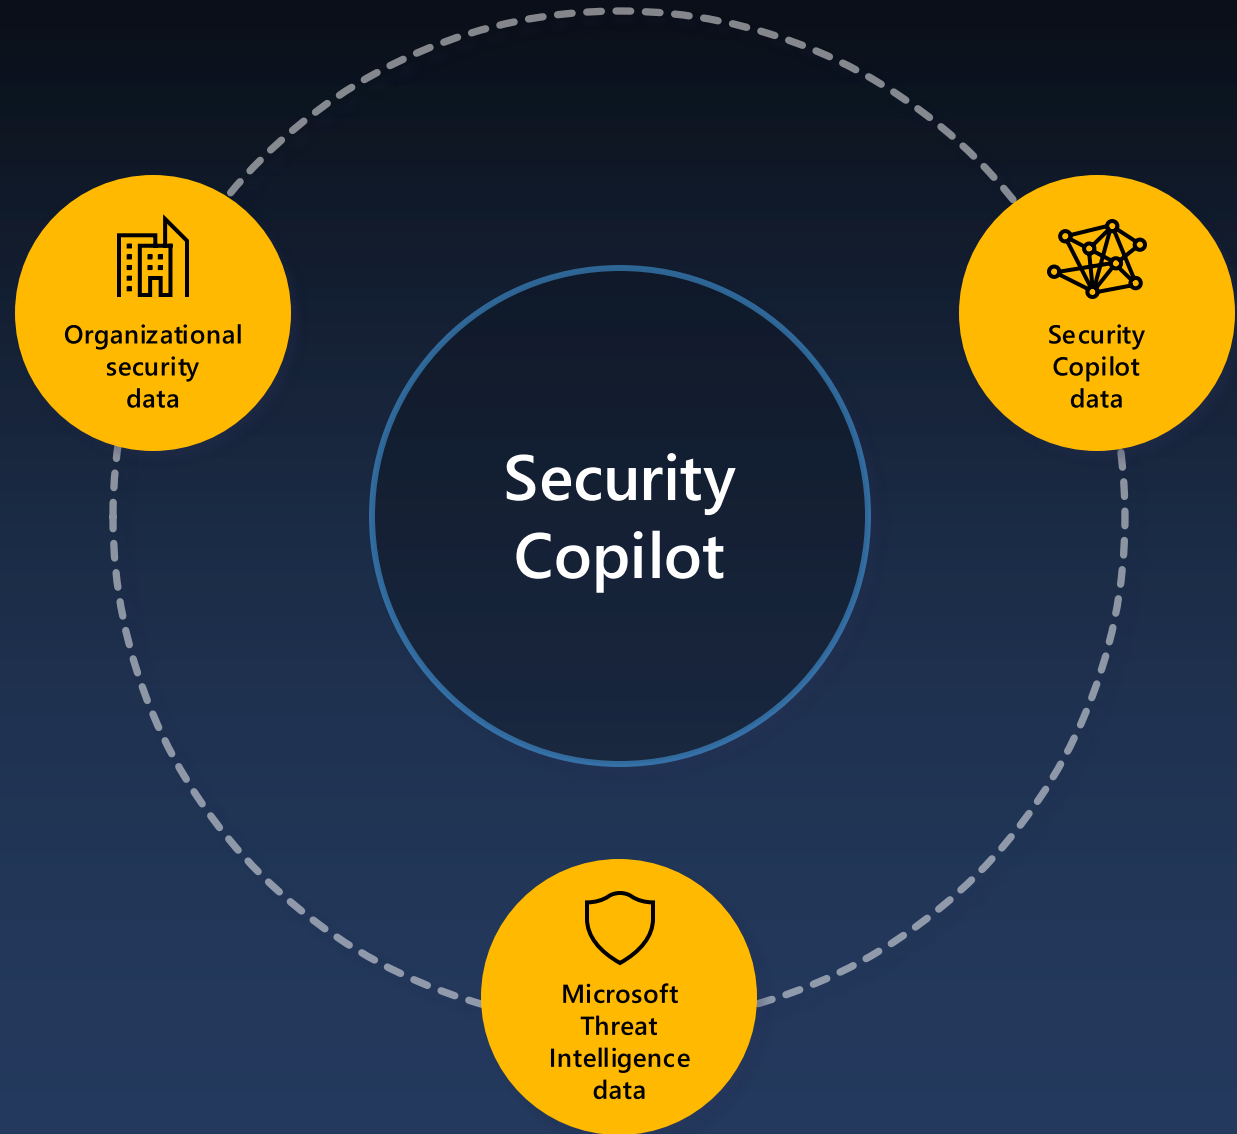
Users spend time automating repeatable and manual tasks to optimize their workflows. Despite efforts, these tasks are traditionally personal and not always broadly shared across an organization.

We created a concept called Promptbooks that are a set of prompts that run to accomplish a specific workflow. Individuals or organizations can build and publish their own or leverage one from the broader community.

## **Impact**

1. No longer need to know a skill in order to do the work
2. Changes the way we get work done
3. Users can learn by using
4. Community building and possible revenue generation

Powered by data  
that is unique to  
you and your  
organization.



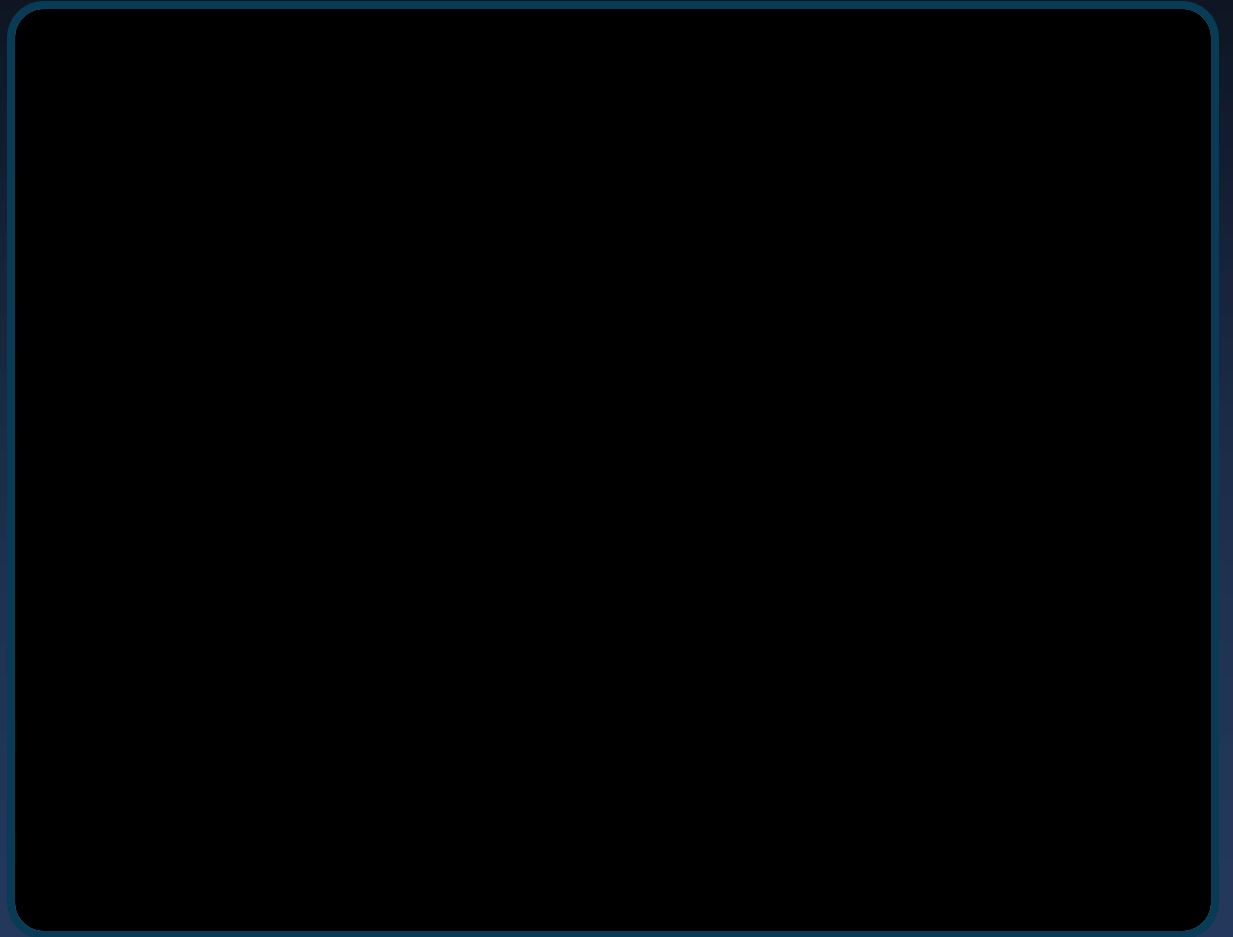
# Going beyond thumbs up/down

With large language model (LLM), the feedback loop is not just supplemental, it is core to the development of the model. It is important to consider various methods in which feedback is obtained. Thumbs up and thumbs down doesn't meet the longer-term needs for training the model.

We're actively exploring new approaches to make feedback more embedded in core interactions.

## **Impact**

1. Engage and empower a user to provide feedback
2. Design both implicit and explicit interactions
3. Create affordances which measure quality
4. Create affordances which infer accountability
5. Collect robust telemetry to measure and improve



Built with security,  
privacy, and  
compliance.

Your data is **your** data



Your data is **not** used to train  
the foundation AI models



Your data is protected by the  
**most comprehensive** enterprise  
compliance and security controls



# Security Copilot working with Microsoft Security



## Microsoft Defender for Endpoint

Monitor devices in real-time  
Detect and prevent threats  
Control policy and access  
Respond to incidents and hunt



## Microsoft Sentinel

Manage logs  
Detect advanced threats  
Monitor and alert in real-time  
Get compliance and reporting



## Microsoft Intune

Manage device inventory  
Enforce configurations and policies  
Deploy and update software  
Deliver conditional access

Security Copilot

- Run queries using natural language
- Prepare reports, summaries, and graphs
- Upskill teams via prompts and guidance
- Reverse engineer malware
- Enrich alerts



- Run queries using natural language
- Prepare reports, summaries, and graphs
- Upskill teams via prompts and guidance
- Reverse engineer malware
- Enrich alerts
- Enrich incidents




- Run queries using natural language
- Prepare reports, summaries, and graphs
- Upskill teams via prompts and guidance
- Reverse engineer malware
- Enrich alerts
- Enrich incidents
- Assess security posture of devices



# Sila Gen AI – sumár incidentu

Popíš mi Incident 2516 Ako  
ROZPRÁVKU pre 3 ročné dieťa 😊

can you describe Defender Incident 2516 as a fairytale for 3 years old kids?

>  3 steps completed 13 seconds

Oct 7, 5:39 PM

Once upon a time, in a magical kingdom called Cyberland, there was a brave knight named Sir Defender. Sir Defender's job was to protect the kingdom from sneaky dragons and mischievous trolls who wanted to cause trouble.

One sunny day, Sir Defender received a message from the wise old owl, who kept an eye on the kingdom from the tallest tree. The message said, "Incident 2516! A sneaky dragon is trying to steal the kingdom's treasure!"

Sir Defender quickly put on his shiny armor and rode his trusty horse, Sparkle, to the treasure vault. When he arrived, he saw the sneaky dragon trying to break in. The dragon was using a magic spell to unlock the door.

But Sir Defender was clever and brave. He knew just what to do! He took out his magical shield, which could block any spell, and held it up high. The dragon's spell bounced off the shield and disappeared into the sky.

The dragon, seeing that it couldn't win, flew away with a roar, promising never to return. The kingdom's treasure was safe once again, thanks to Sir Defender.

The people of Cyberland cheered and celebrated, and Sir Defender was hailed as a hero. And from that day on, everyone in the kingdom slept soundly, knowing that Sir Defender was always there to protect them from any danger.

And they all lived happily ever after. The end.

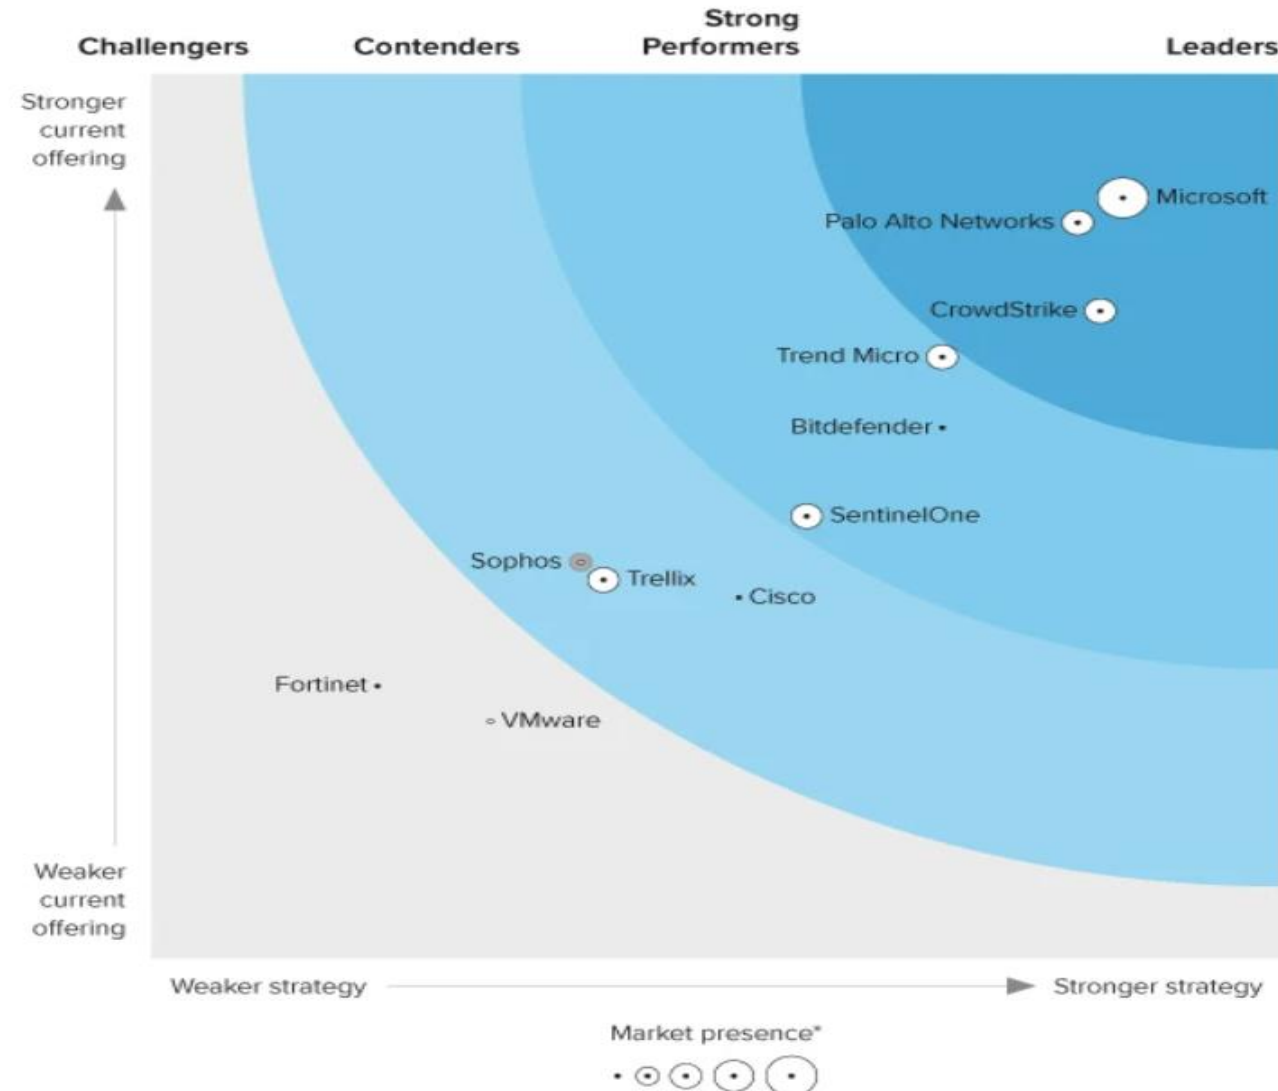
😊 How's this response?



# THE FORRESTER WAVE™

## Extended Detection And Response Platforms

Q2 2024



\*A gray bubble or open dot indicates a nonparticipating vendor.

Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.

Ďakujem za pozornosť!





Can you summarize sentinel incident 29088 contained in workspace woodgrove-loganalyticsworkspace?

> 3 steps completed 1 min



2:09 PM

Incident 29088 in the Woodgrove-LogAnalyticsWorkspace has the title "SAP - (Preview) File Downloaded From a Malicious IP Address" and is described as a user downloading a file from an SAP system using an IP address known to be malicious. The incident has a severity of "High" and is currently in "New" status. The incident was first detected on 2023-06-21T18:38:41.103Z and last updated on 2023-06-22T18:43:44.4744196Z.

The incident has 1 associated alert with the same title and description. The alert was generated on 2023-06-22T18:43:44.1397144Z and has a tactic of "Exfiltration".

There are 3 entities associated with this incident:

1. IP: 185.82.217.3
2. Host: CPC-mscot-S0C0S (Windows, FQDN: cpc-mscot-s0c0s)
3. Account: [mscott@woodgrove.ms](#) (AzureAD domain, DisplayName: [mscott@woodgrove.ms](#))

Please note that this is a summary of the available information, and further investigation may be required to fully understand the incident.

Confirm

Off-target

Report

Pin

Ask anything about security, or type / for suggestions



- Last External IP Address: 20.172.15.250

- Last Internal IP Address: 10.1.0.9

- AV Status: NotSupported

### 3. Account: [mscott@woodgrove.ms](#)

- Account Name: mscott

- NT Domain: AzureAD

- UPN Suffix: woodgrove.ms

- SID: S-1-12-1-3834083884-1079035900-3132596361-2383408245

- AAD Tenant ID: 536279f6-15cc-45f2-be2d-61e352b51eef

- AAD User ID: e4877a2c-c7fc-4050-89a0-b7ba75ec0f8e

- Is Domain Joined: true

- Display Name: [mscott@woodgrove.ms](#)

- Department: Sales

- Job Title: Sales II

- Office Location: Studio-X/101

- City: Redmond

- Company Name: Woodgrove

- Country: US

- State: WA

- Street Address: 1 Microsoft Way

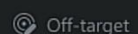
- User Principal Name: [mscott@woodgrove.ms](#)

- Additional Mail Addresses: [mscott@hipstabank.com](#)

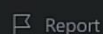
These entities are involved in the incident where a user downloaded a file from an SAP system using a malicious IP address.



Confirm



Off-target



Report



Pin



Ask anything about security, or type / for suggestions



Confirm Off-target Report Pin



You note the IP address is malicious, can you provide the reputation and more details about why it's malicious?

> 3 steps completed 1 min



2:13 PM

The IP address 185.82.217.3 has a reputation score of 100, which classifies it as malicious. The following rules contribute to its malicious classification:

**1. Cyber Threat Intelligence - Cobalt Strike**

- Description: Cobalt Strike
- Severity: 5
- Link: [Cobalt Strike](#)

**2. Cyber Threat Intelligence - Silk Typhoon**

- Description: Silk Typhoon
- Severity: 5
- Link: [Silk Typhoon](#)

**3. ASN**

- Description: Infrastructure hosted by this ASN frequently exhibits suspicious behavior
- Severity: 4
- Link: None

These rules indicate that the IP address is associated with known cyber threat intelligence profiles and suspicious behavior.

Confirm Off-target Report Pin



Ask anything about security, or type / for suggestions





• Link: None

These rules indicate that the IP address is associated with known cyber threat intelligence profiles and suspicious behavior.



Confirm



Off-target



Report



Pin



Tell me more about cobalt strike and silk typhoon.



3 steps completed 1 min



+2

2:14 PM

Here's more information about Cobalt Strike and Silk Typhoon:

#### 1. Cobalt Strike

Cobalt Strike is a commercial, full-featured penetration testing tool that allows an attacker to deploy an agent named "Beacon" on the victim's machine. Beacon includes a wealth of functionality to the attacker, including command execution, keylogging, file transfer, SOCKS proxying, privilege escalation, and more. It is often used by advanced persistent threat (APT) groups and cybercriminals for post-exploitation activities. Cobalt Strike has been observed in various cyber attacks, including those taking advantage of the Exchange Server vulnerabilities. You can read more about it in this [article](#).

#### 2. Silk Typhoon (HAFNIUM)

Silk Typhoon is a threat actor group that has been observed exploiting the Exchange Server vulnerabilities. They are known for their sophisticated attacks and have targeted various organizations worldwide. Microsoft researchers have observed a steady increase in signals indicating exploitation by threat actors other than Silk Typhoon. Numerous other attackers have been seen adopting the exploit into their toolkits. You can read more about Silk Typhoon and their activities in this [article](#).



Confirm



Off-target



Report



Pin



Ask anything about security, or type / for suggestions








```
1 let userEmails = dynamic(["mscott@woodgrove.ms", "mscott@hipstabank.com"]);
2 union
3 (
4     AlertEvidence
5     | where AccountUpn in (userEmails) or AccountName in (userEmails)
6     | extend userEmail = iif(isnull(AccountUpn), AccountName, AccountUpn)
7 ),
8 (
9     EmailEvents
10    | where SenderFromAddress in (userEmails) or RecipientEmailAddress in (userEmails)
11    | extend userEmail = iif(SenderFromAddress in (userEmails), SenderFromAddress,
12                           RecipientEmailAddress)
13 )
14 | where Timestamp > datetime("2023-07-12T17:57:34Z") - 30d
15 | project Timestamp, userEmail, Title, AlertId, NetworkMessageId, EmailDirection
```

I found a total of 15 records, and here's the table with the results:

15 rows × 6 columns



Timestamp	UserEmail	Title
2023-06-21T23:07:31Z	mscott@woodgrove.ms	
2023-06-28T23:36:52Z	mscott@woodgrove.ms	
2023-06-28T22:48:19Z	mscott@woodgrove.ms	
2023-07-06T23:14:39Z	mscott@woodgrove.ms	
2023-06-23T00:00:50Z	mscott@woodgrove.ms	

 Confirm  Off-target  Report


 Pin

 Ask anything about security, or type / for suggestions





these devices compliant.

>  3 steps completed 1 min



2:15 PM

Here are the devices associated with the user [mscott@woodgrove.ms](#) and their compliance status:

3 rows × 6 columns



Device Name	Manufacturer	Model
MSCOTT-SURFACE2	Microsoft Corporation	Surface Laptop 2
DESKTOP-LDLSMJL	Microsoft Corporation	Virtual Machine
CPC-mscot-S0C0S	Microsoft Corporation	Cloud PC Enterprise 2vCPU/8GB/128GB

To make these devices compliant, follow these steps:

**1. Remain in contact**


Ensure that the devices remain in contact with the Intune service. Devices that do not remain in contact may be considered noncompliant. Check the device's network connectivity and ensure it can communicate with the Intune service.


**2. Defender for Endpoint - Device Risk**

For the device CPC-mscot-S0C0S, ensure that the required security level for Device Threat Protection is met. This can be achieved by updating the device's security software, applying security patches, and ensuring that the device is protected against known threats.


After addressing these issues, the devices should become compliant. Keep in mind that it may take some time for the compliance status to update in Intune.

 Confirm

 Off-target

 Report

 Pin

 Ask anything about security, or type / for suggestions



The entities involved in this incident are:

- Host: CPC-mscot-S0C0S (Windows)
- IP: 185.82.217.3
- Account: [mscott@woodgrove.ms](#) (AzureAD)

### Supporting Evidence

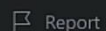
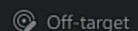
The IP address 185.82.217.3 has a reputation score of 100, classifying it as malicious. It is associated with known cyber threat intelligence profiles, such as Cobalt Strike and Silk Typhoon, and has exhibited suspicious behavior.


### Confidence Level

Based on the available evidence, the confidence level in the assessment of this incident is high.

### Recommendations

1. Investigate the user's activities and the file downloaded from the malicious IP address to determine if any sensitive data was exfiltrated or if the user's account was compromised.
2. Review the security policies and access controls for the user and their devices to ensure they are in line with the organization's security requirements.
3. Implement additional monitoring and alerting for the user's account and associated devices to detect any potential malicious activities in the future.
4. Educate the user on the risks associated with downloading files from untrusted sources and provide guidance on how to identify and avoid such threats.
5. Ensure that all devices associated with the user are compliant with the organization's security policies and have up-to-date security software and patches installed.



 Ask anything about security, or type / for suggestions