



Trellix

SOC 3.0

Automate Investigations with AI

Nate Mayer
Principal Solution Architect
July 10, 2024

Presentation Overview

History of Machine Learning and GenAI at Trellix

Development of Trellix Wise

GenAI Auto-Investigations

Trellix Wise with EDR + Helix

Trellix Wise + SOAR

37 Year Heritage


founded

FireEye acquired
MANDIANT

Trellix
founded

 
expanded
integrations

1987 1996 1998 2004 2007 2011 2013 2017 2020 2021 JAN 2022 MAR 2022 SEP 2022 FEB 2023 APR 2023 MAY 2023 JUL 2023 SEP 2023 DEC 2023 FEB 2024


founded


founded


IDC names Trellix as
Modern Endpoint
Security Leader

Rich History of Machine Learning & Artificial Intelligence

2014
Created analytics such
as Impossible Travel

2016
Launched Guided
investigations

2016
Launched NextGen EPP
with ML

2023
Launched AI driven auto
investigations

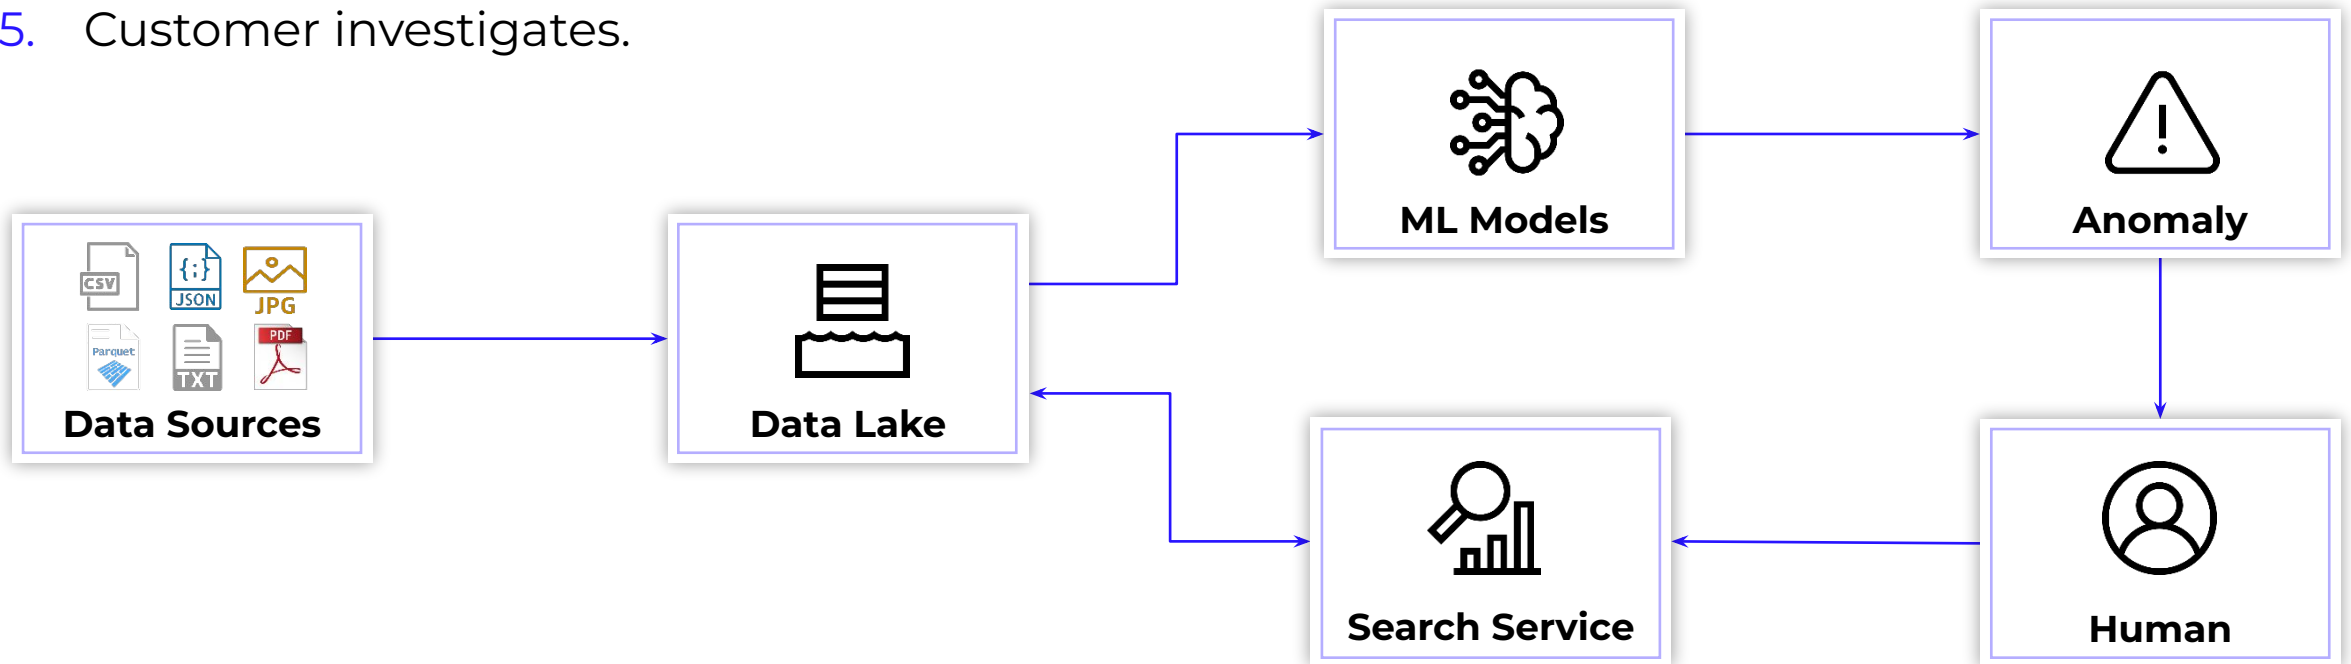


Trellix

How Trellix Uses Generative AI

Our pre-AI approach

1. Create thousands of connectors and parsers to normalize event data from anywhere.
2. Store all of the data on S3 and OpenSearch.
3. Analyze the data for anomalies with Amazon EMR and ML models.
4. Report findings back to the customer.
5. Customer investigates.



Effective, but hard to scale

What Worked

Ingesting Data
Analyzing and Matching
Searching obscene data volumes

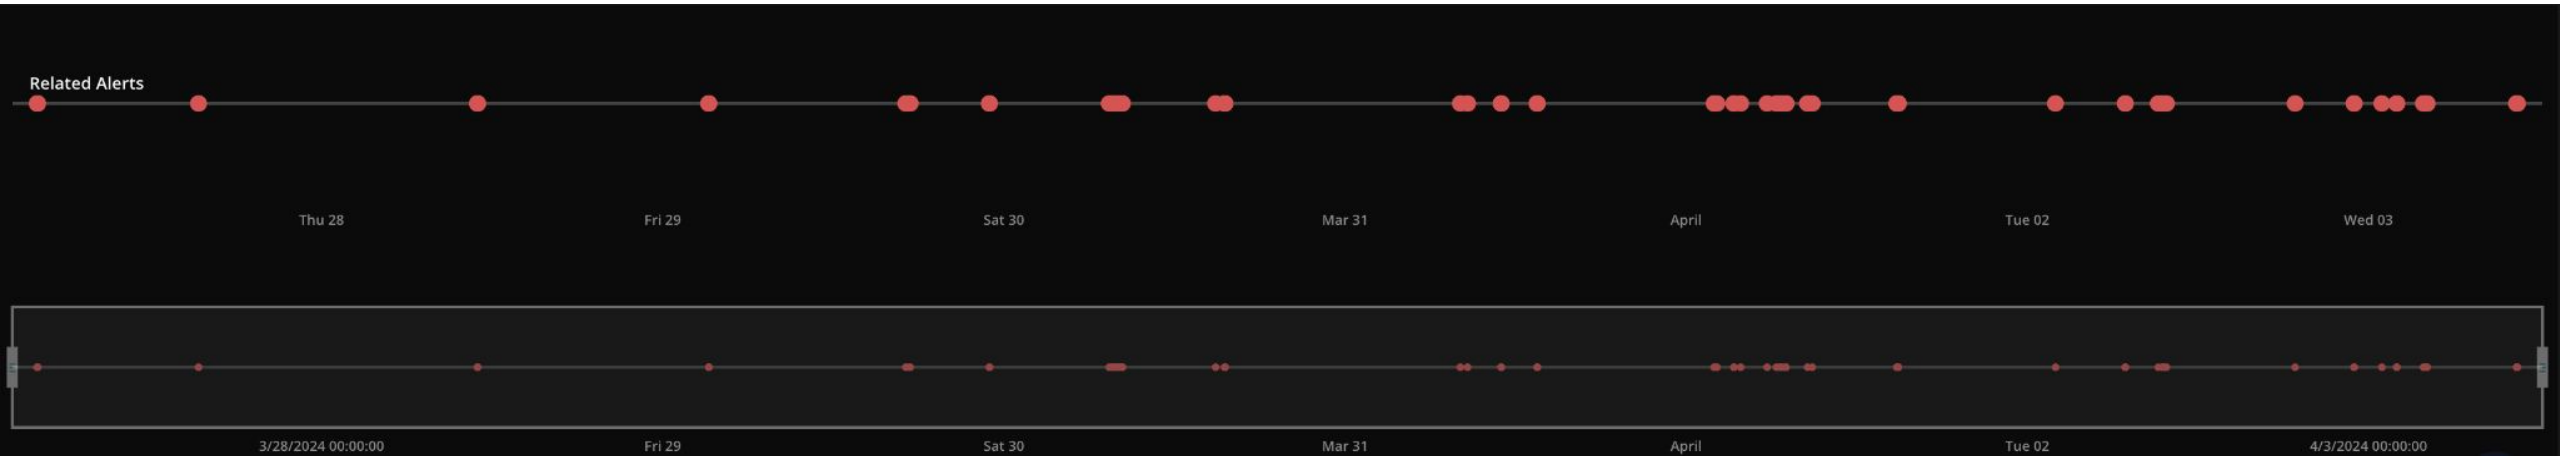
What Didn't

Didn't have time to investigate all findings



Example Security Timeline

Analytics timeline



Abnormal AWS AMI usage, API key usage, key pairs generated

Abnormal console, app, and directory logins

Abnormal office suite, directory, and infrastructure commands

How do we find time to investigate everything that is “weird?”

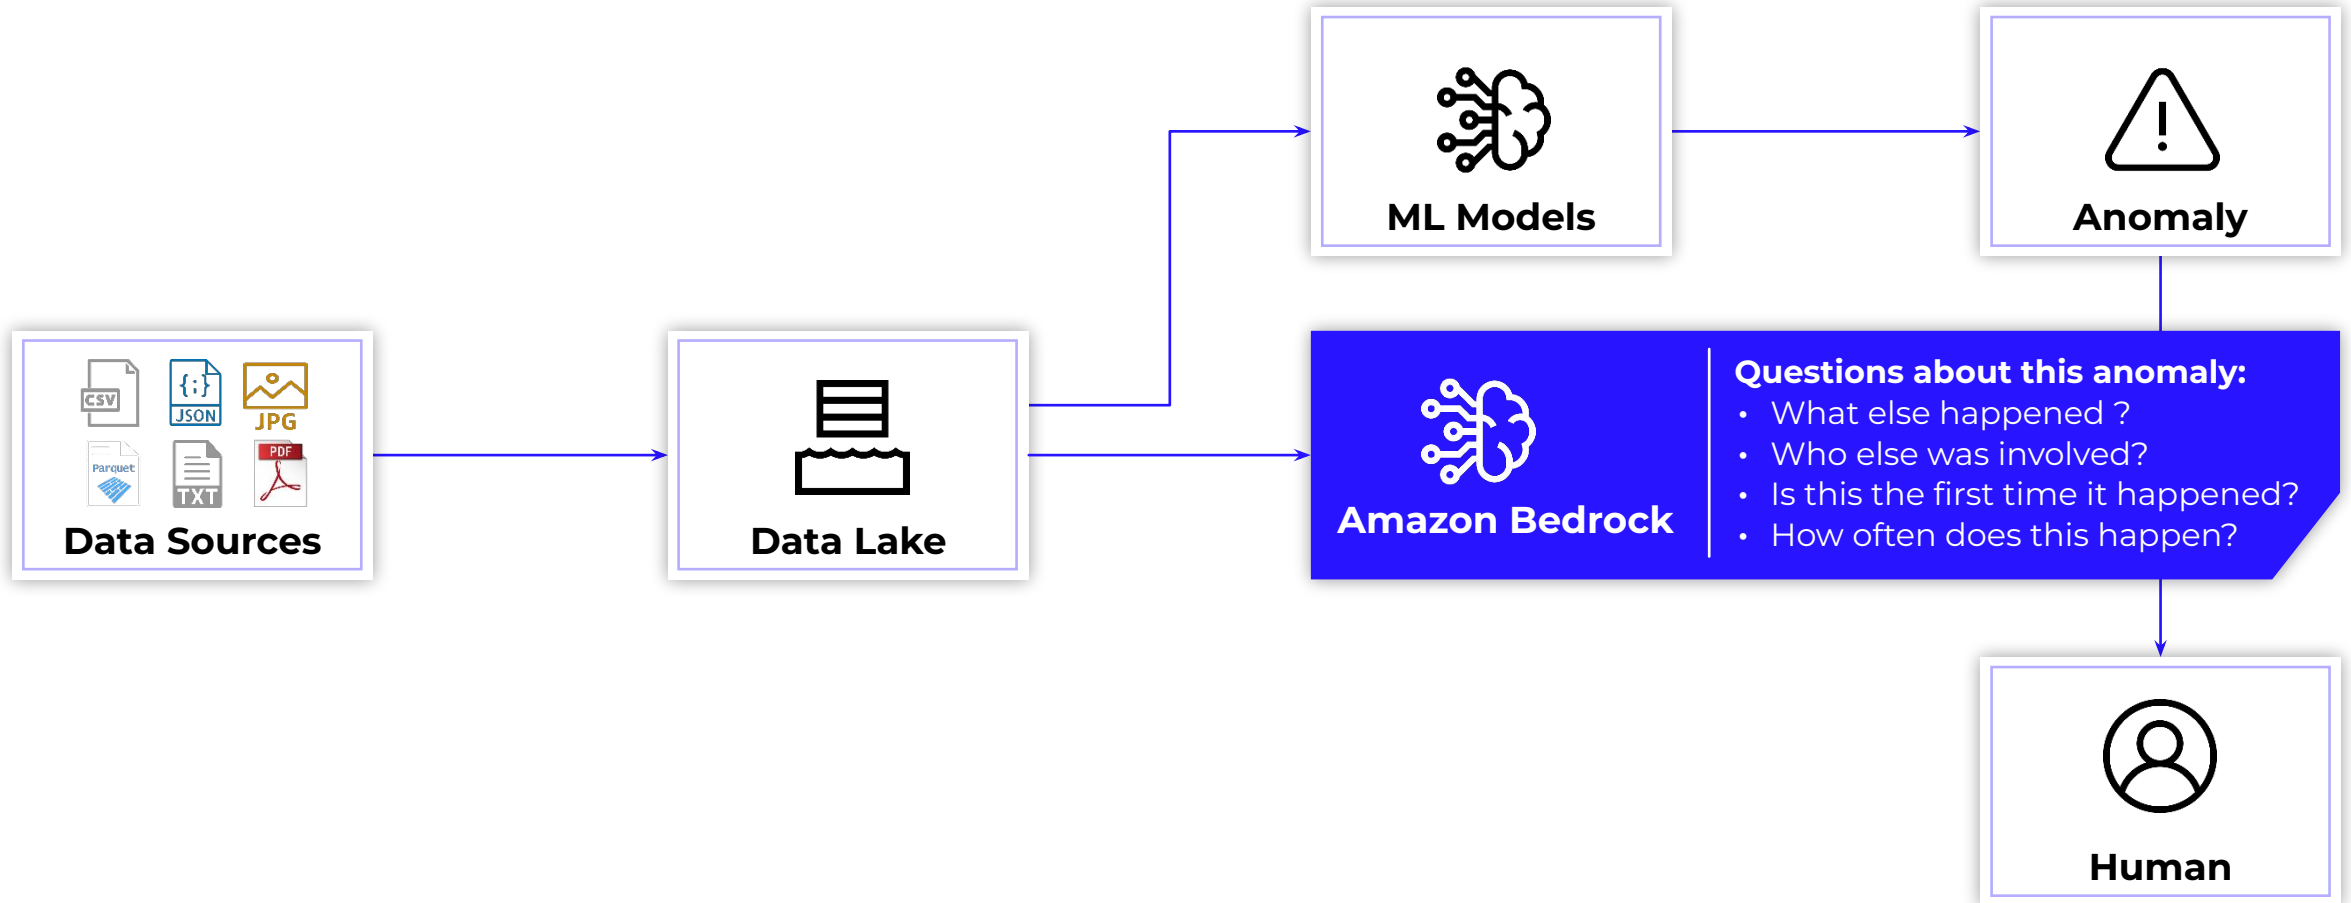


Detecting anomalies was not enough.

We needed to focus on the right signals.

We needed gen AI.

Generative AI can ask key questions and understand answers



What it takes to make gen AI work

1. Initial findings to investigate
2. Sub-second data retrieval times for all answers
3. Pre-built investigations for generative AI to ask the right questions



Find me the Top Ten Things I need to look at



Example: Anomalous Console Login

Analytic thesis:



Based on recent investigations, the tell-tale sign that an account is compromised is a password reset followed by a login from a new country.

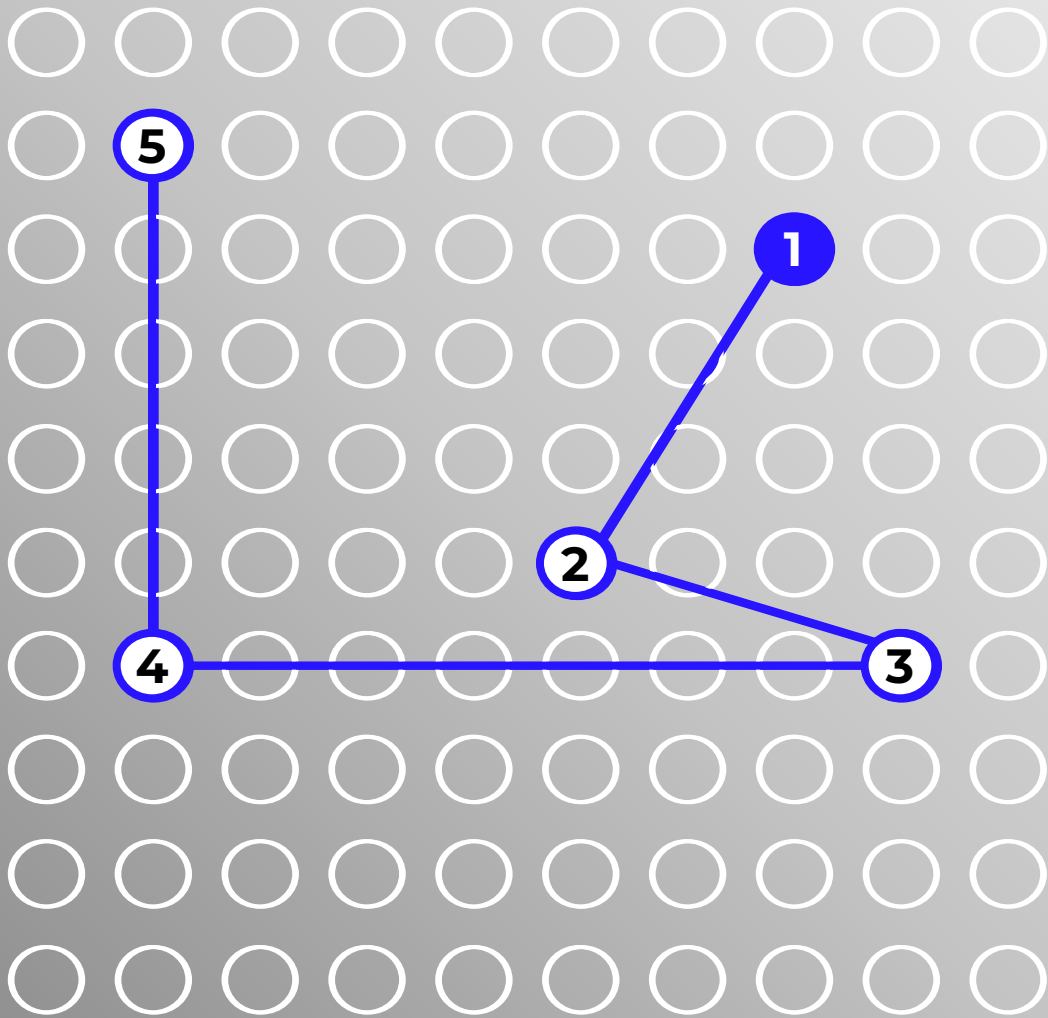
Therefore, if a user resets a password and logs in from somewhere atypical, that needs to be investigated.

Real World:



Users travel, forget passwords, have executive assistants, etc.

This happens hundreds of times per day in large organizations.



Generative AI Solution:

1. Get some answers to specific questions (human time saved):
2. Is this user currently traveling? (3m)
3. Were there failed attempts before the password reset? (1m)
4. Does the user have an executive assistant? (1m)
5. What level of access does this user have? (3m)
6. What did they do after they reset the password? (5m)

Generative AI performs **13 minutes** of work in under **45 seconds**.

How does GenAI actually interpret investigation results?

Based on the information provided, there are several signs of suspicious activity from IP 10.207.17.85 that **warrant raising the alert level**:

- The SIEM alert indicates a detected brute force attack correlated with successful logins from the same IP. This **alone is a serious security event**.
- The IP has **triggered other rules recently** including Trellix IPS rules, indicating ongoing malicious behavior.
- The IP has generated logs across various systems including firewalls, DNS, and Windows events, showing **broad suspicious activity**.

Given the brute force attack, other security events, and signs of compromise across multiple systems, **I would recommend raising the alert level for this event to High severity**. A brute force attack that results in a successful login is a critical security incident that requires priority investigation and response. The **supporting evidence** of other suspicious network activity from this IP makes the situation more severe.



Trellix

Real-world examples of GenAI auto-investigations

Investigate alert: TRELIX ENDPOINT HX

[file write to named pipe (methodology)]

GenAI ANALYSIS:

Based on the information provided, there are a few factors to consider:

- The IOC detected relates to use of named pipes for command and control, which is inherently suspicious activity.
- The same IOC was triggered by 2 other hosts in the last 24 hours.
- The source IP has communicated with several external IPs, some of which **appear to be cloud hosting providers**.
- The source hostname and user are **tied to other recent alerts**, indicating repeated suspicious activity.
- The affected asset is the **'system' account, which has high privileges** on the host.

While none of these are definitive signs of compromise on their own, together they paint a picture of potentially malicious activity that warrants further investigation. The source system and user should be examined more closely through additional logging and monitoring.

I would recommend raising the severity of this alert to High, to **escalate the response**. The clustering of related activity and criticality of the affected asset justify treating this with increased urgency. However, more conclusive evidence would likely be needed to declare this a verified incident. Continued monitoring and gathering of forensic data should be prioritized.

Investigate alert: TRELIX ENDPOINT HX

[fe_trojan_macro_generic_129]

GenAI ANALYSIS:

Based on the information provided, there are a few signs of suspicious activity that warrant raising the alert level:

- The Trellix HX alert indicates detection of a file write IOC on the affected asset. This is a potential sign of malware infection.
- The supporting rules show additional Trellix HX alerts for malware detection and IOCs on the same source IP. This strengthens the indication of potential compromise.
- There are intel hits related to this source IP for non-DNS FQDNs associated with malware. This is further corroboration.
- The affected asset triggered the same IOC detection for a Trojan macro. The repeated detection is concerning.
- The MD5 hash for this alert is **associated with other alerts for this asset and IP**. This indicates the malware is being spread.
- The user associated with **the asset is a demo user**, which can be more prone to compromise if credentials are shared.

Given the above, I would recommend raising the alert level to High severity. While a single IOC detection may only warrant a Medium alert, the repeated detections, intel hits, and signs of malware spread justify escalating the severity. Prompt investigation and containment actions should be prioritized for this alert.

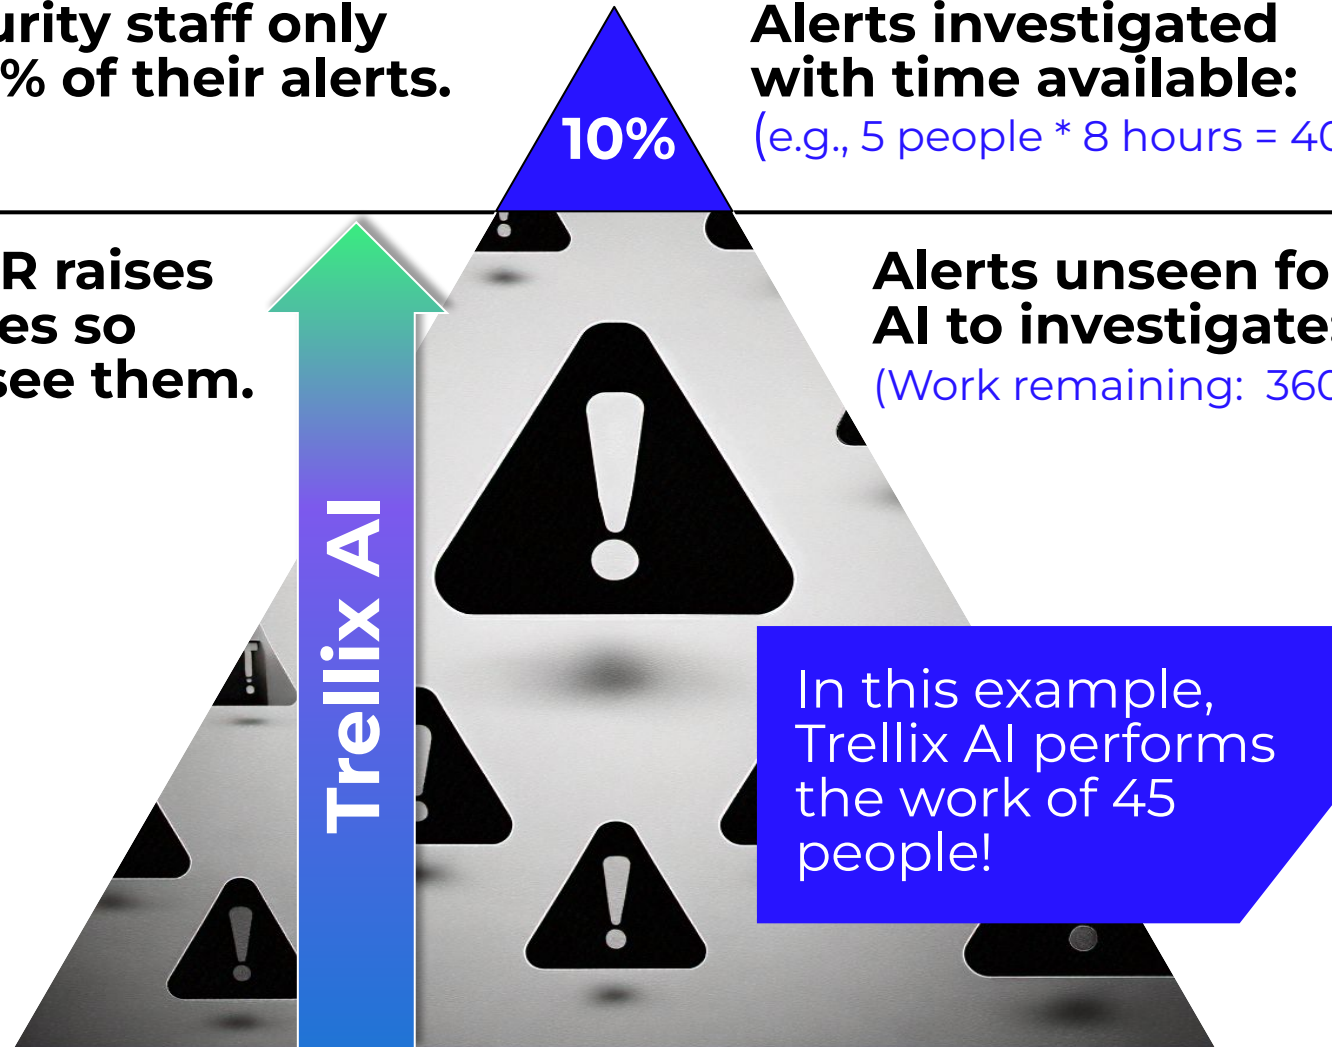
Never miss an alert

Most security staff only look at 10% of their alerts.

Alerts investigated with time available:
(e.g., 5 people * 8 hours = 40 hours/day)

Trellix XDR raises alert scores so analysts see them.

Alerts unseen for AI to investigate:
(Work remaining: 360 hours/day)



Finally, there is enough help to get the job done.



Trellix scales
Analysts
with auto
investigations



Average customer scenario

Alerts Per day	1,100
Events Per Alert	65
Analyst time per event	5 Seconds

To analyze every alert, Bedrock performs **357,500** seconds of analyst time per day, which is about **12** 8-hour shifts.

Tuning across the entire ecosystem

Helix Connect allows Trellix Wise to be given specific instructions and guidance for its decision making.

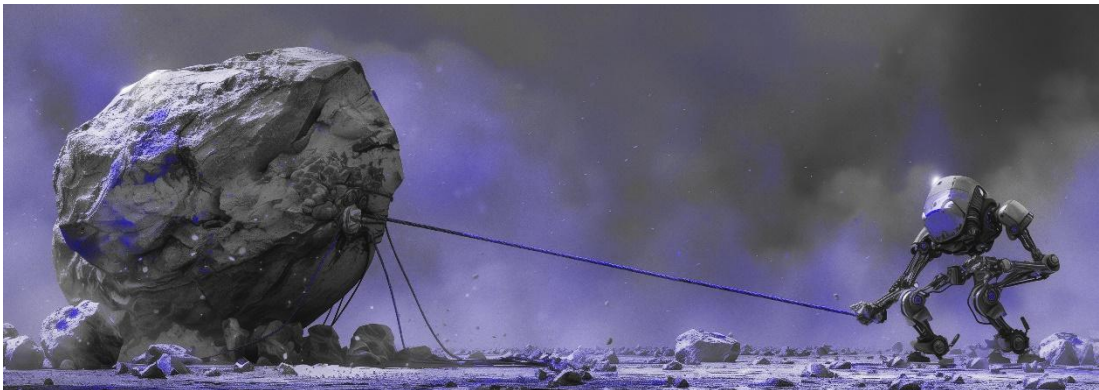
This can be anything. Examples:

Always escalate endpoint alerts when the user has access to AWS.

Only escalate alerts from endpoints belonging to sales on weekends.

Be more suspicious of phishing emails near the end of the fiscal quarter.

Evolve from data mining to alert mining



Before

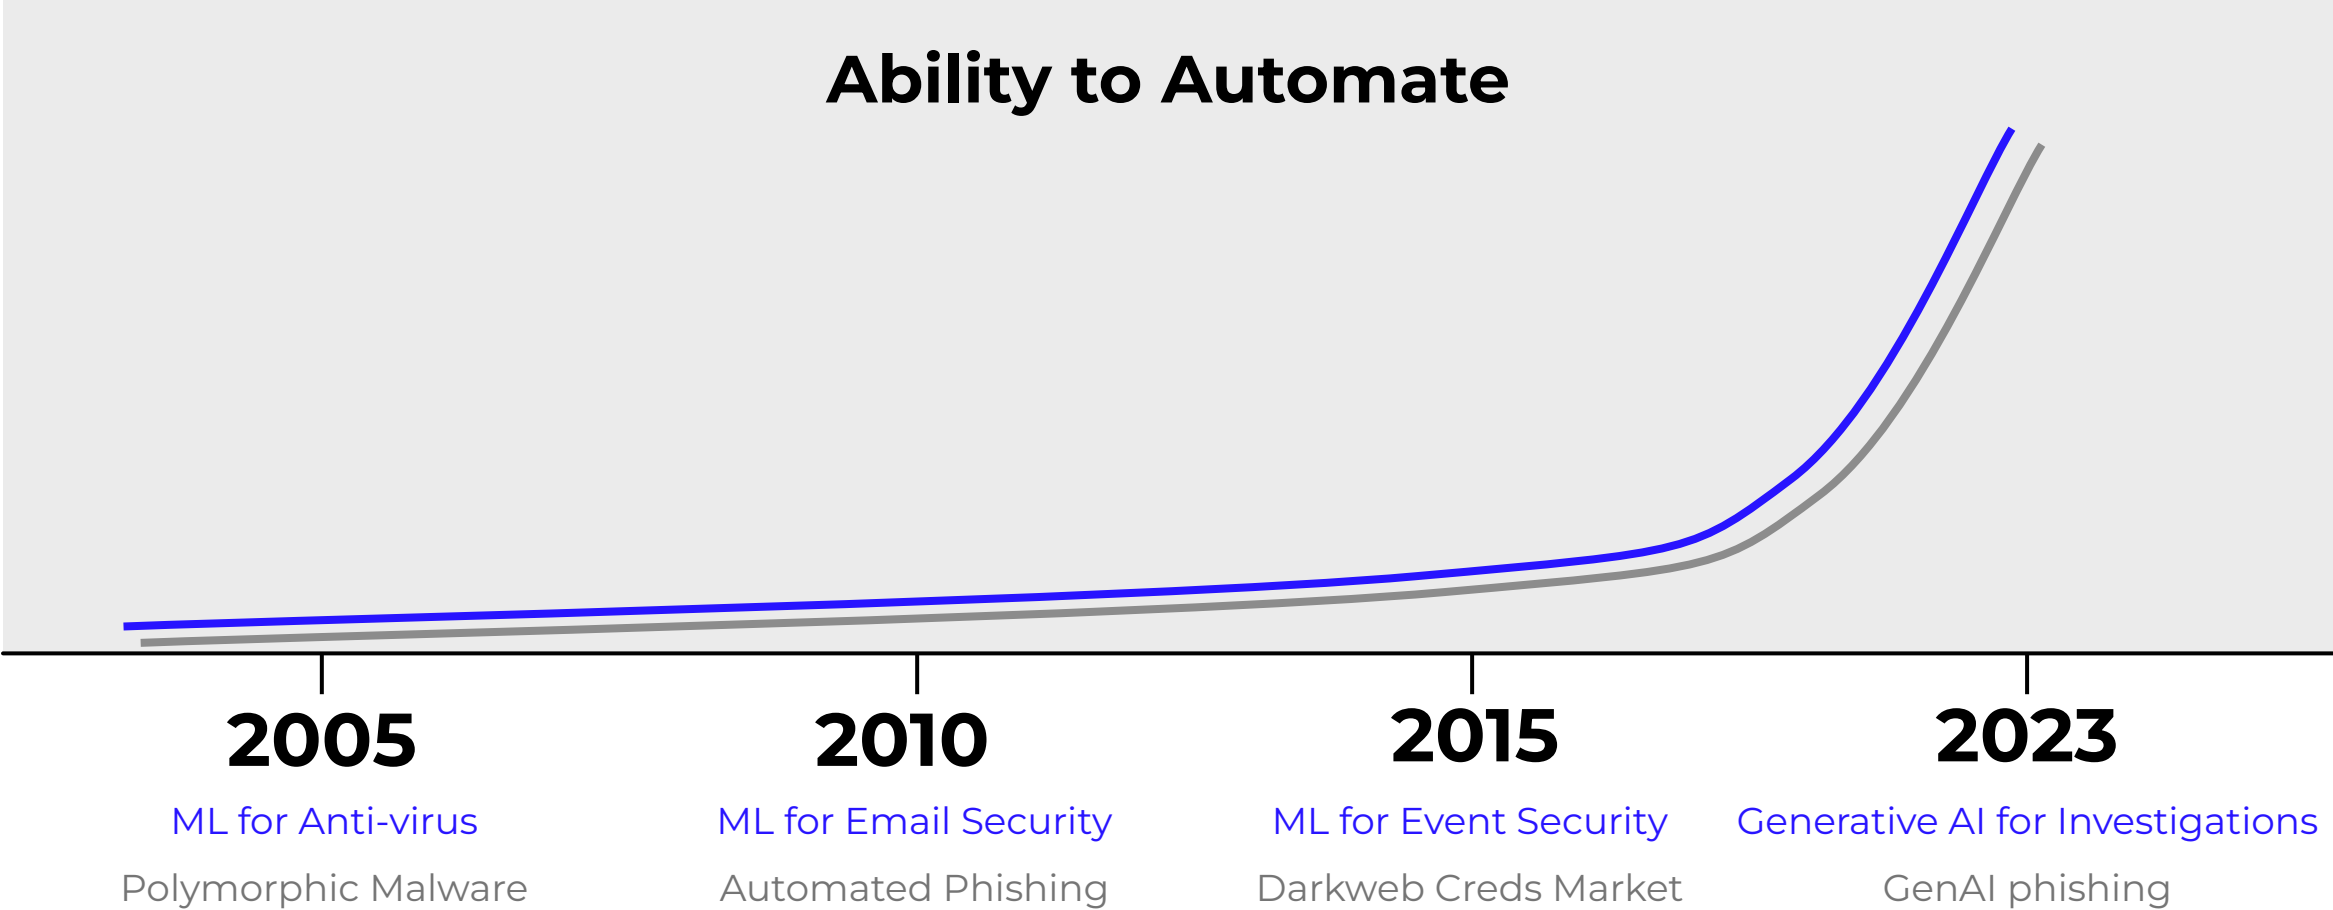
- Analyst overwhelmed by alerts
- Waste time tuning tools to reduce alerts
- Only investigate alerts that are clear/obvious
- Reduce alert aperture to known-bad
- Ignore most alerts



After

- Focus on top 1% without penalty
- Turn on all available alert sources
- Deep investigations on most valuable alerts
- Spend time on innovation and threat hunting
- No alerts ignored

The AI arms race





Trellix

Trellix Wise with EDR

Trellix Wise with EDR

Use Cases

- Natural language query for Historical and Real Time Search
- Rightsized Security Posture Management
- Accelerated investigations and threat hunting
- Dossier Mode provides executive summaries of an incident
- Interactive Mode enables analysts to uncover new security insights
- Knowledge Graph visually shows the attack path



Natural Language Search

Trellix Wise with EDR

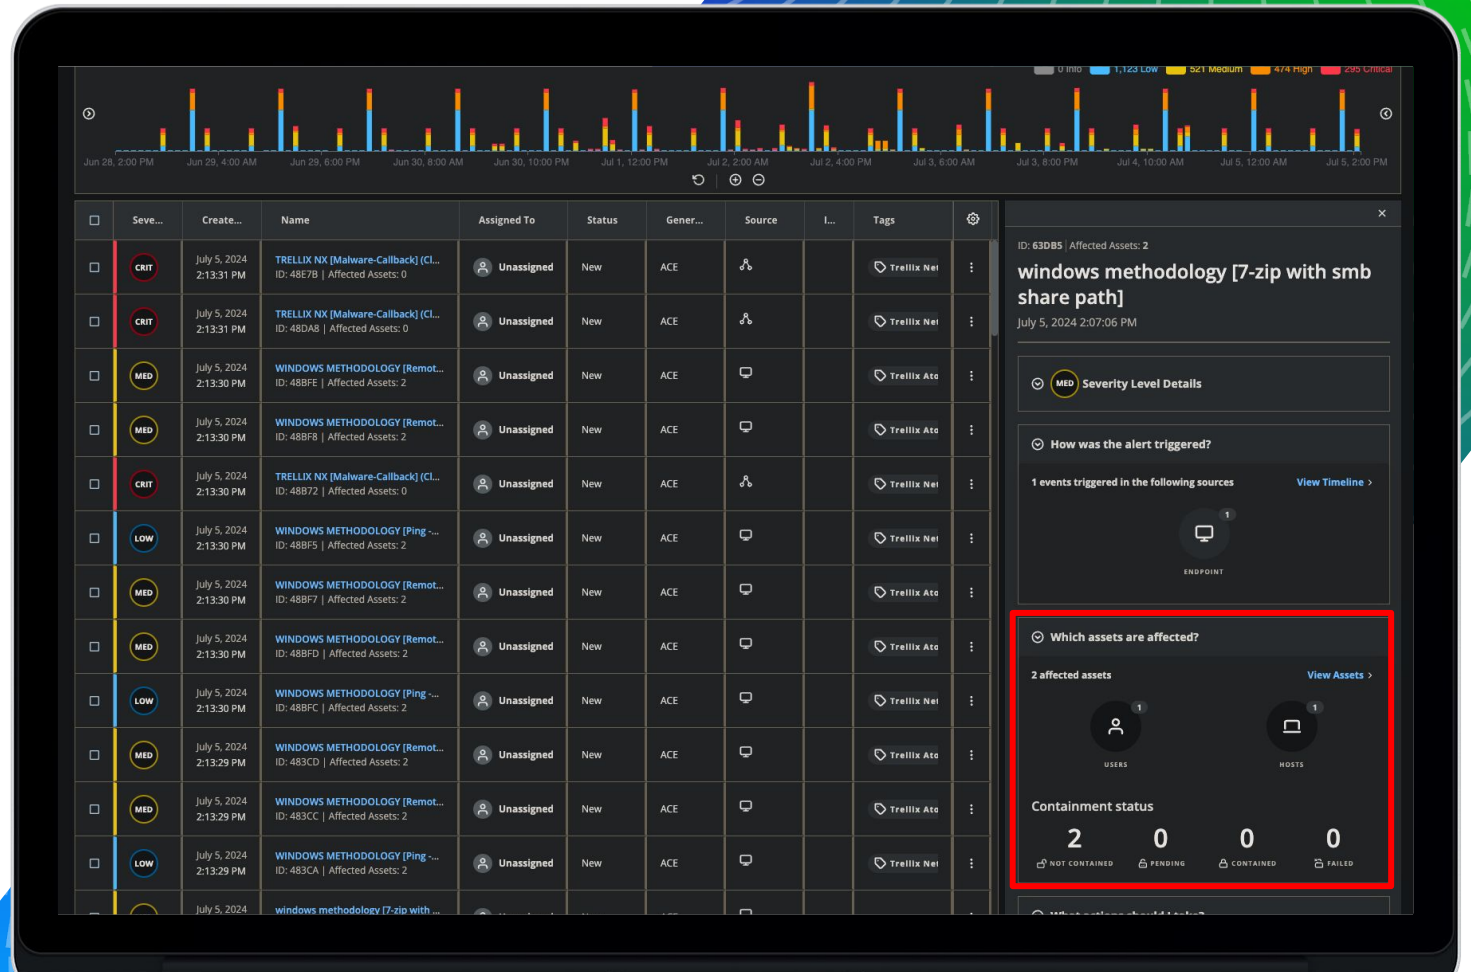
The screenshot shows the Trellix EDR Historical Search interface. At the top left, there is a Trellix logo and a navigation path: EDR → Historical Search. Below this, the 'Historical Search' section is active. A search bar contains the text 'Search with Wise' and a dropdown arrow. To the right of the search bar, a green checkmark icon is followed by the natural language query: 'show me all events for 192.168.10.1 and for the endpoint named x-laptop'. A search icon and a filter icon labeled 'Last 4 hours' are also visible. Below the search bar, a box labeled 'GENERATED QUERY' displays the resulting query: 'IpAddress = "192.168.10.1" AND DeviceName = "x-laptop"'. The interface has a dark theme.

Use natural language to determine all the events pertaining to an endpoint, file, or process. For example, you can ask:

“show me all events for 192.168.10.1 and for the endpoint named x-laptop”

Analyze Detection – Affected Assets

Quickly view assets generating alerts and their status – Not Contained, Pending, Contained



Analyze Detection – Suggested Actions

Quickly view GenAI-suggested actions to be performed based on alert generated and asset type

The screenshot displays a security dashboard interface. At the top, there are filters for 'Time Range: Last 7 Days', 'Assignee: Me, Unassigned', 'Status: New, In Progress', and 'Generated By: All'. A bar chart shows alert counts over time, with a legend for severity levels: Info (0), Low (1,123), Medium (521), High (474), and Critical (295). Below the chart is a table of alerts with columns for Severity, Create time, Name, Assigned To, Status, Generation, Source, Location, and Tags. A detailed view of an alert is shown on the right, titled 'WINDOWS METHODOLOGY [Remote AT Usage]' with ID 48BF8 and 2 affected assets. The alert was triggered on July 5, 2024, at 2:13:31 PM. The 'How was the alert triggered?' section shows 1 event triggered in the following sources, with a link to 'View Timeline'. The 'Which assets are affected?' section is also visible. A dashed blue box highlights the 'What actions should I take?' section, which lists recommended actions: 'Check if the flagged activity is legitimate or expected', 'Run an initial analysis to determine severity, impact, and attack flow', 'Stop the infected process', and 'Delete malicious file, scheduled task, and service'.

Seve...	Create...	Name	Assigned To	Status	Gener...	Source	L...	Tags
CRIT	July 5, 2024 2:13:31 PM	TRELLIX NX [Malware-Callback] (CL... ID: 48E7B Affected Assets: 0	Unassigned	New	ACE			Trellix Nei
CRIT	July 5, 2024 2:13:31 PM	TRELLIX NX [Malware-Callback] (CL... ID: 48DA8 Affected Assets: 0	Unassigned	New	ACE			Trellix Nei
MED	July 5, 2024 2:13:30 PM	WINDOWS METHODOLOGY [Remot... ID: 48BF8 Affected Assets: 2	Unassigned	New	ACE			Trellix Ato
MED	July 5, 2024 2:13:30 PM	WINDOWS METHODOLOGY [Remot... ID: 48BF8 Affected Assets: 2	Unassigned	New	ACE			Trellix Ato
CRIT	July 5, 2024 2:13:30 PM	TRELLIX NX [Malware-Callback] (CL... ID: 48B72 Affected Assets: 0	Unassigned	New	ACE			Trellix Nei
LOW	July 5, 2024 2:13:30 PM	WINDOWS METHODOLOGY [Ping ... ID: 48BF5 Affected Assets: 2	Unassigned	New	ACE			Trellix Nei
MED	July 5, 2024 2:13:30 PM	WINDOWS METHODOLOGY [Remot... ID: 48BF7 Affected Assets: 2	Unassigned	New	ACE			Trellix Ato
MED	July 5, 2024 2:13:30 PM	WINDOWS METHODOLOGY [Remot... ID: 48BFD Affected Assets: 2	Unassigned	New	ACE			Trellix Ato
LOW	July 5, 2024 2:13:30 PM	WINDOWS METHODOLOGY [Ping ... ID: 48BFC Affected Assets: 2	Unassigned	New	ACE			Trellix Nei
MED	July 5, 2024 2:13:29 PM	WINDOWS METHODOLOGY [Remot... ID: 48BCD Affected Assets: 2	Unas	New	ACE			Trellix Ato
MED	July 5, 2024 2:13:29 PM	WINDOWS METHODOLOGY [Remot... ID: 48BCC Affected Assets: 2	Unas	New	ACE			Trellix Ato

What actions should I take?

Recommended actions

- Check if the flagged activity is legitimate or expected
- Run an initial analysis to determine severity, impact, and attack flow
- Stop the infected process
- Delete malicious file, scheduled task, and service

Trellix Detection – Modes

Easily switch between Dossier mode for executive summaries and Interactive mode to unearth new insights through guided threat hunting

The screenshot displays the Trellix EDR Monitoring dashboard. At the top, it shows 'Monitoring' with a summary of 18 Total Threats, categorized into 9 High, 3 Medium, and 6 Low. A 'Threats by Ranking' list on the left includes items like 'code-injection-sample.exe' and 'Artemis-VeryLow-Applicatio...'. The main panel shows details for a 'code-injection-sa...' threat, including its initial trigger, detection times, affected devices, and process attributes. A 'Threat Behavior' section lists techniques observed, such as 'MITRE ATT&CK™' and 'OS Credential Dumping T1003'. A 'Process Activity' section shows a summary view of the threat's activity, including file paths like 'powershell.exe'. On the right, a 'Wise' AI summary panel is open, offering to generate an executive summary. A yellow arrow points from the 'Threat Behavior' section to the 'Wise' panel. A blue callout box at the bottom right contains the text 'Exec Summary'.

Monitoring 18 Total Threats 9 High 3 Medium 6 Low

Threats by Ranking

code-injection-sample.exe May 23, 2024 12:50:37 AM

Artemis-VeryLow-Applicatio... May 29, 2024 7:46:07 AM

Artemis-VeryLow.exe May 29, 2024 4:05:21 AM

Artemis-VeryLow-Troj.exe May 29, 2024 4:05:21 AM

Threat-Sample2_1WNVr.exe May 13, 2024 3:10:44 AM

Threat-Sample.exe May 28, 2024 1:47:15 AM

code-injection-sa...
Initial trigger
First detection May 13, 2024 1:25:32 AM
Last detection May 23, 2024 12:50:37 AM
Affected devices 4
Age 22 days
Take Action

Process Attributes
First Name code-injection-sample.exe
MD5 E2E43E9D0C479E0E5D41BCB4EA9FA1B4
SHA-1 28E9D7E4158FCFBAC28A9187390AB7918CE08F44
SHA-256 EAA94C9E35D2EE3C003A4F636D5B67DF9644BB97D54325CD43E93FE10CB95146

Threat Details
Device: 671W10R55X64 May 13, 2024 1:25:32 AM
Threat Behavior
Techniques Observed(49) MITRE ATT&CK™
OS Credential Dumping T1003 (Credential Access)
LSASS Memory T1003.001 (Credential Access)
Security Account Manager T1003.002 (Credential Access)
NTDS T1003.003 (Credential Access)
LSA Secrets T1003.004 (Credential Access)
Process Activity
Summary View
code-injection-sample.exe
code-injection-sample.exe powershell.exe image path \\ccda...ode-i...

Wise
*Generated by AI, verify for accuracy
Welcome to Trellix's new AI capabilities! Choose an option to test it out.
Select Mode:
 Interactive mode
Chat bot like interaction with initial summary plus suggested contextual prompts
 Dossier mode
One off document analysis report
Triage

Exec Summary

Analyze Detection – Dossier Mode

Dossier mode provides executive summaries of an incident that details what happened, where it happened, when it happened, and whose credentials were involved.

The screenshot displays the Trellix EDR Monitoring interface in Dossier Mode. The top navigation bar shows 'Trellix | EDR → Monitoring' and a status summary: 18 Total Threats, 9 High, 3 Medium, and 6 Low. The main content area is divided into several sections:

- Threats by Ranking:** A list of threats with a search filter and a 'View' dropdown set to 'All'. The list includes 'code-injection-sample.exe', 'Artemis-VeryLow-Applicatio...', 'Artemis-VeryLow.exe', 'Artemis-VeryLow-Troj.exe', 'Threat-Sample2_1WNVr.exe', 'Threat-Sample.exe', and 'Keep-Running.exe'.
- Threat Details:** A detailed view for 'code-injection-sample.exe'. It shows:
 - Initial trigger:** code-injection-sample.exe
 - Trace detection:** May 13, 2024 1:25:32 AM
 - First detection:** May 13, 2024 1:25:32 AM
 - Last detection:** May 23, 2024 12:50:37 AM
 - Affected devices:** 4
 - Age:** 22 days
 - Process Attributes:** First Name: code-injection-sample.exe; MD5: E2E43E9D0C479E0E5D41BCB4E9FA1B4; SHA-1: 28E9D7E4158FCFBAC28A9187390AB7918CE08F44; SHA-256: EAA94C9E35D2EE3C003A4F636D5B67DF9644BB97D54325CD43E93FE10CB95146
 - Threat Behavior:** Techniques Observed (49) including OS Credential Dumping T1003 (Credential Access), LSASS Memory T1003.001 (Credential Access), Security Account Manager T1003.002 (Credential Access), NTDS T1003.003 (Credential Access), and LSA Secrets T1003.004 (Credential Access).
 - Process Activity:** A table showing the execution of 'code-injection-sample.exe', 'powershell.exe', and 'whoami.exe'. The 'code-injection-sample.exe' entry is marked as 'Observed' and includes an 'Image path' of '\\cdafilesrv001.beaeng.mfe.exe' and a 'Parent process'.
- Detection Analysis:** A summary of the events indicating a potential credential access and privilege escalation attack. It mentions the execution of the 'Invoke-Mimikatz PowerShell script' and provides key points for investigation.

Analyze Detection – Interactive Mode

Interactive mode enables the unearthing of new insights and their MITRE mappings through guided threat hunting by helping analysts answer questions of

When did the incident happen?

What do I do with this information?

What actions can I take?

Where can I get more information?

The screenshot displays the Trellix EDR Monitoring dashboard. At the top, it shows 18 Total Threats, with 9 High, 3 Medium, and 6 Low severity threats. The main view is 'Monitoring' with a list of threats on the left and a detailed view of a selected threat on the right. The selected threat is 'code-injection-sample.exe', detected on May 23, 2024, at 12:50:37 AM. The interface shows various attributes including Initial trigger, Trace detection, First and Last detection times, Affected devices (4), and Age (22 days). A 'Take Action' button is visible. The 'Threat Details' section shows 'Device: 671W10RS5X64' and '4 affected devices'. Under 'Threat Behavior', it lists techniques observed, such as OS Credential Dumping T1003 (Credential Access), LSASS Memory T1003.001 (Credential Access), Security Account Manager T1003.002 (Credential Access), NTDS T1003.003 (Credential Access), and LSA Secrets T1003.004 (Credential Access). The 'Process Activity' section shows a summary view with a table of processes: 'code-injection-sample.exe', 'powershell.exe', and 'whoami.exe'. A yellow arrow points to the 'Type' field for 'powershell.exe'. The 'Detection Analysis' section provides a summary of the events, indicating a potential credential access and privilege escalation attack. Key points include: 'The Invoke-Mimikatz PowerShell script was executed, which is a known technique for credential dumping.', 'Suspicious processes that accessed the LSASS process and performed system information discovery were observed.', 'The events indicate a potential credential access and privilege escalation attack.', 'The investigation should focus on the process that executed the Invoke-Mimikatz script and any other related suspicious activities.', and 'The host where this activity was first detected is 671W10RS5X64.' At the bottom, there are buttons for 'Provide more Detail', 'Brief me on related MITRE TTPs', 'Generate a Knowledge Graph', and 'Suggest some Recommended Actions'.

Analyze Detection – Knowledge Graph

Knowledge Graph provides a visual representation of the anatomy of an attack

The screenshot displays the Trellix EDR Monitoring dashboard. At the top, it shows 18 Total Threats, with 9 High, 3 Medium, and 6 Low severity threats. The main interface is divided into several sections:

- Threats by Ranking:** A list of detected threats, including 'code-injection-sample.exe', 'Artemis-VeryLow-Applicatio...', 'Artemis-VeryLow.exe', 'Artemis-VeryLow-Troj.exe', 'Threat-Sample2_1VNW.exe', 'Threat-Sample.exe', 'Keep-Running.exe', 'Threat-Sample2_1FRME.exe', 'Threat-Sample2_VPWHM.exe', 'watsonthreat22.bat', 'lntpreater.cmd.exe', 'schtasks.exe', and 'PersistentSample_Above60...'.
- Threat Details:** A detailed view of a threat, showing initial trigger, first and last detection times, affected devices, and process attributes. The process is identified as 'Artemis-VeryLow-Application.exe'.
- Process Activity:** A summary view of the process activity, showing the image path, type, parent process, command line, process ID, user account, and start time.
- Knowledge Graph:** A visual representation of the attack anatomy, showing a flow from '396W10R5S64' through 'winlogon.exe' and 'explorer.exe' to 'Artemis-VeryLow-Application.exe'.

Trellix

Trellix Wise with SOAR



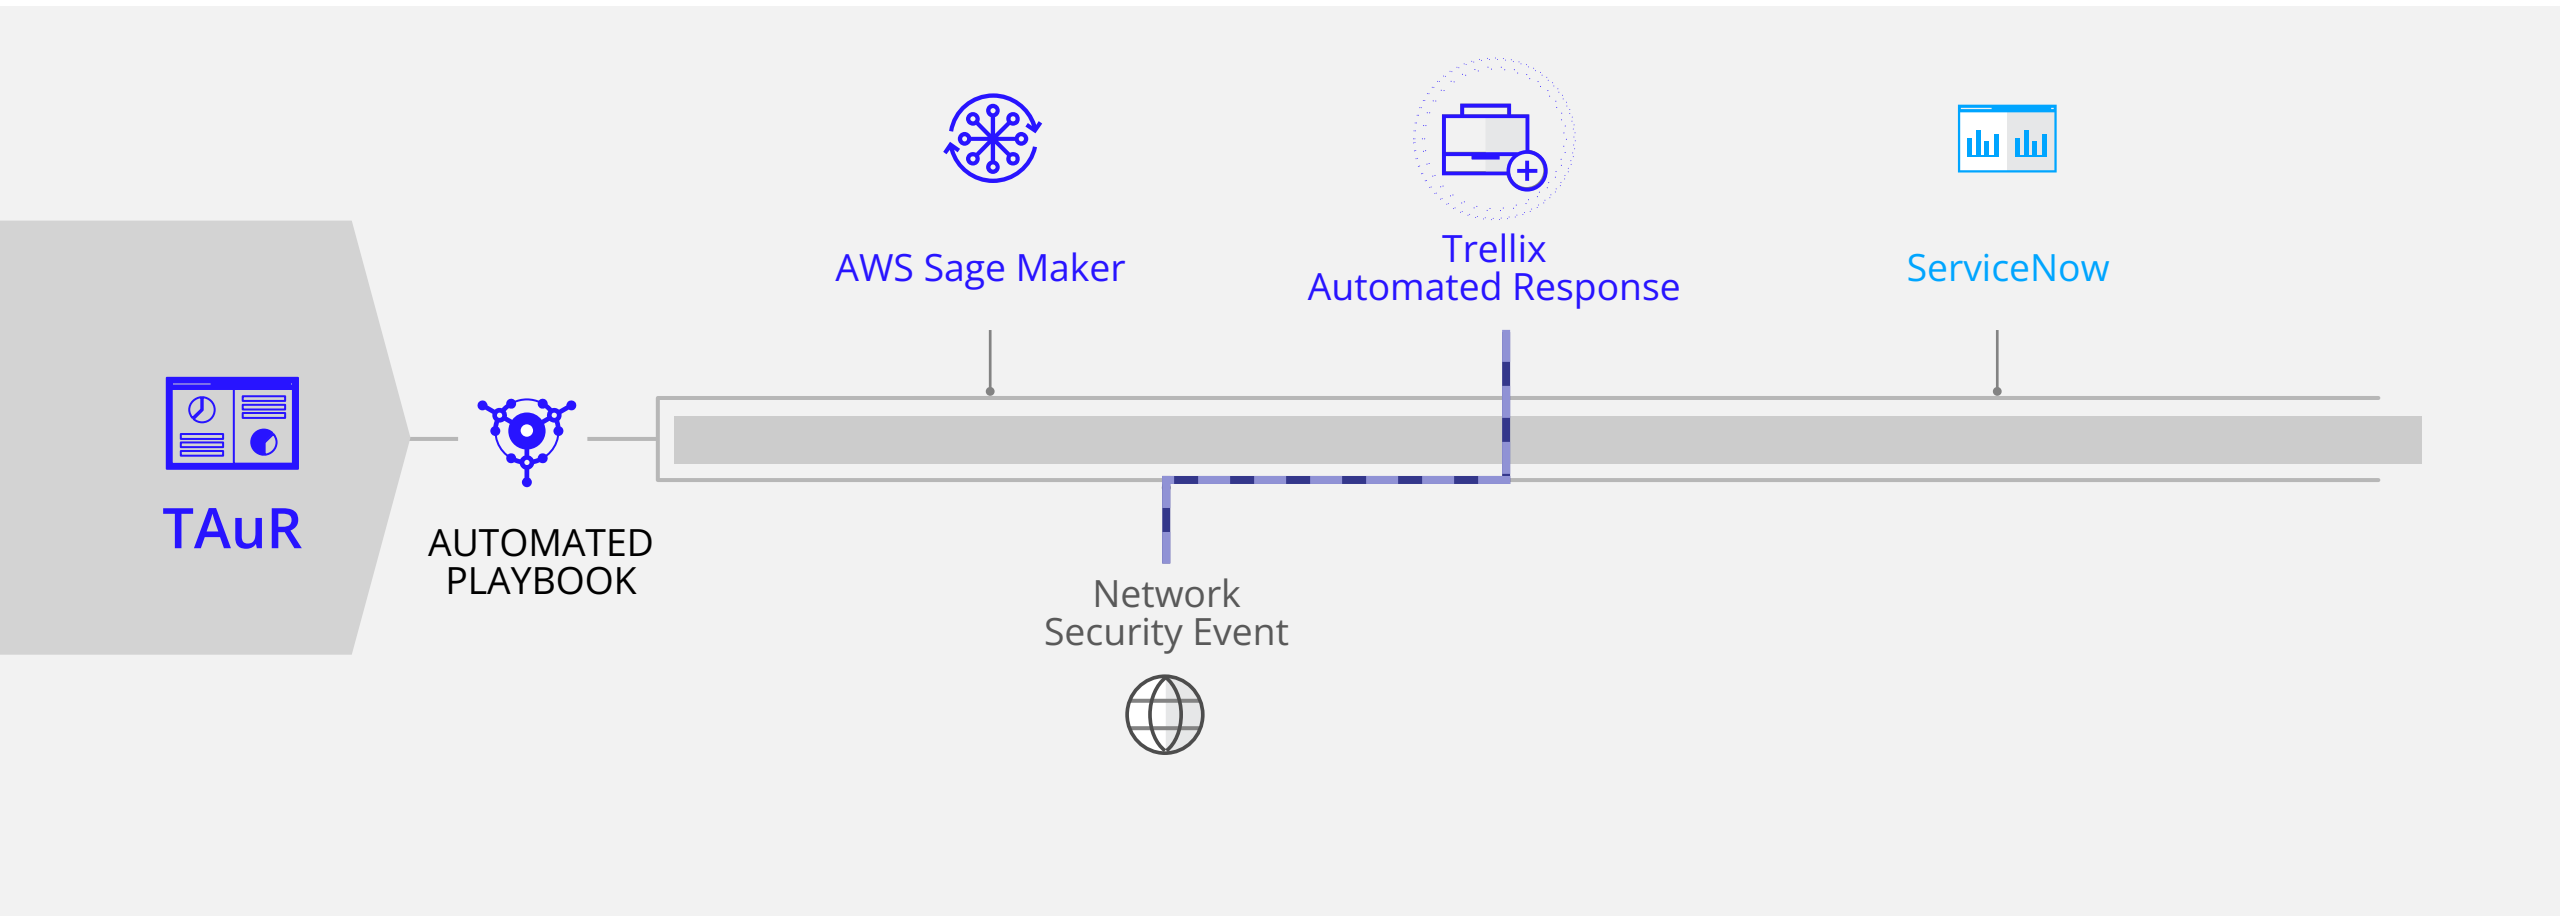
SOAR + GEN AI =

Example Use Cases

- **Summarization of an incident**
 - AI can sift through diverse data sources, including logs, alerts, and threat intelligence, to extract meaningful insights
- **Recommendations for remediation**
 - By analyzing the characteristics and similarities between current and previously resolved incidents, AI can suggest appropriate remediation steps based on proven best practices
- **Recommendations for protections**
 - Based on this analysis, AI can offer recommendations for enhancing defenses, such as implementing intrusion detection and prevention systems, tightening access controls, updating security policies, or conducting security awareness training
- **Multilingual support**
 - AI's language processing capabilities enable SOAR platforms to support multiple languages, overcoming language barriers in incident response

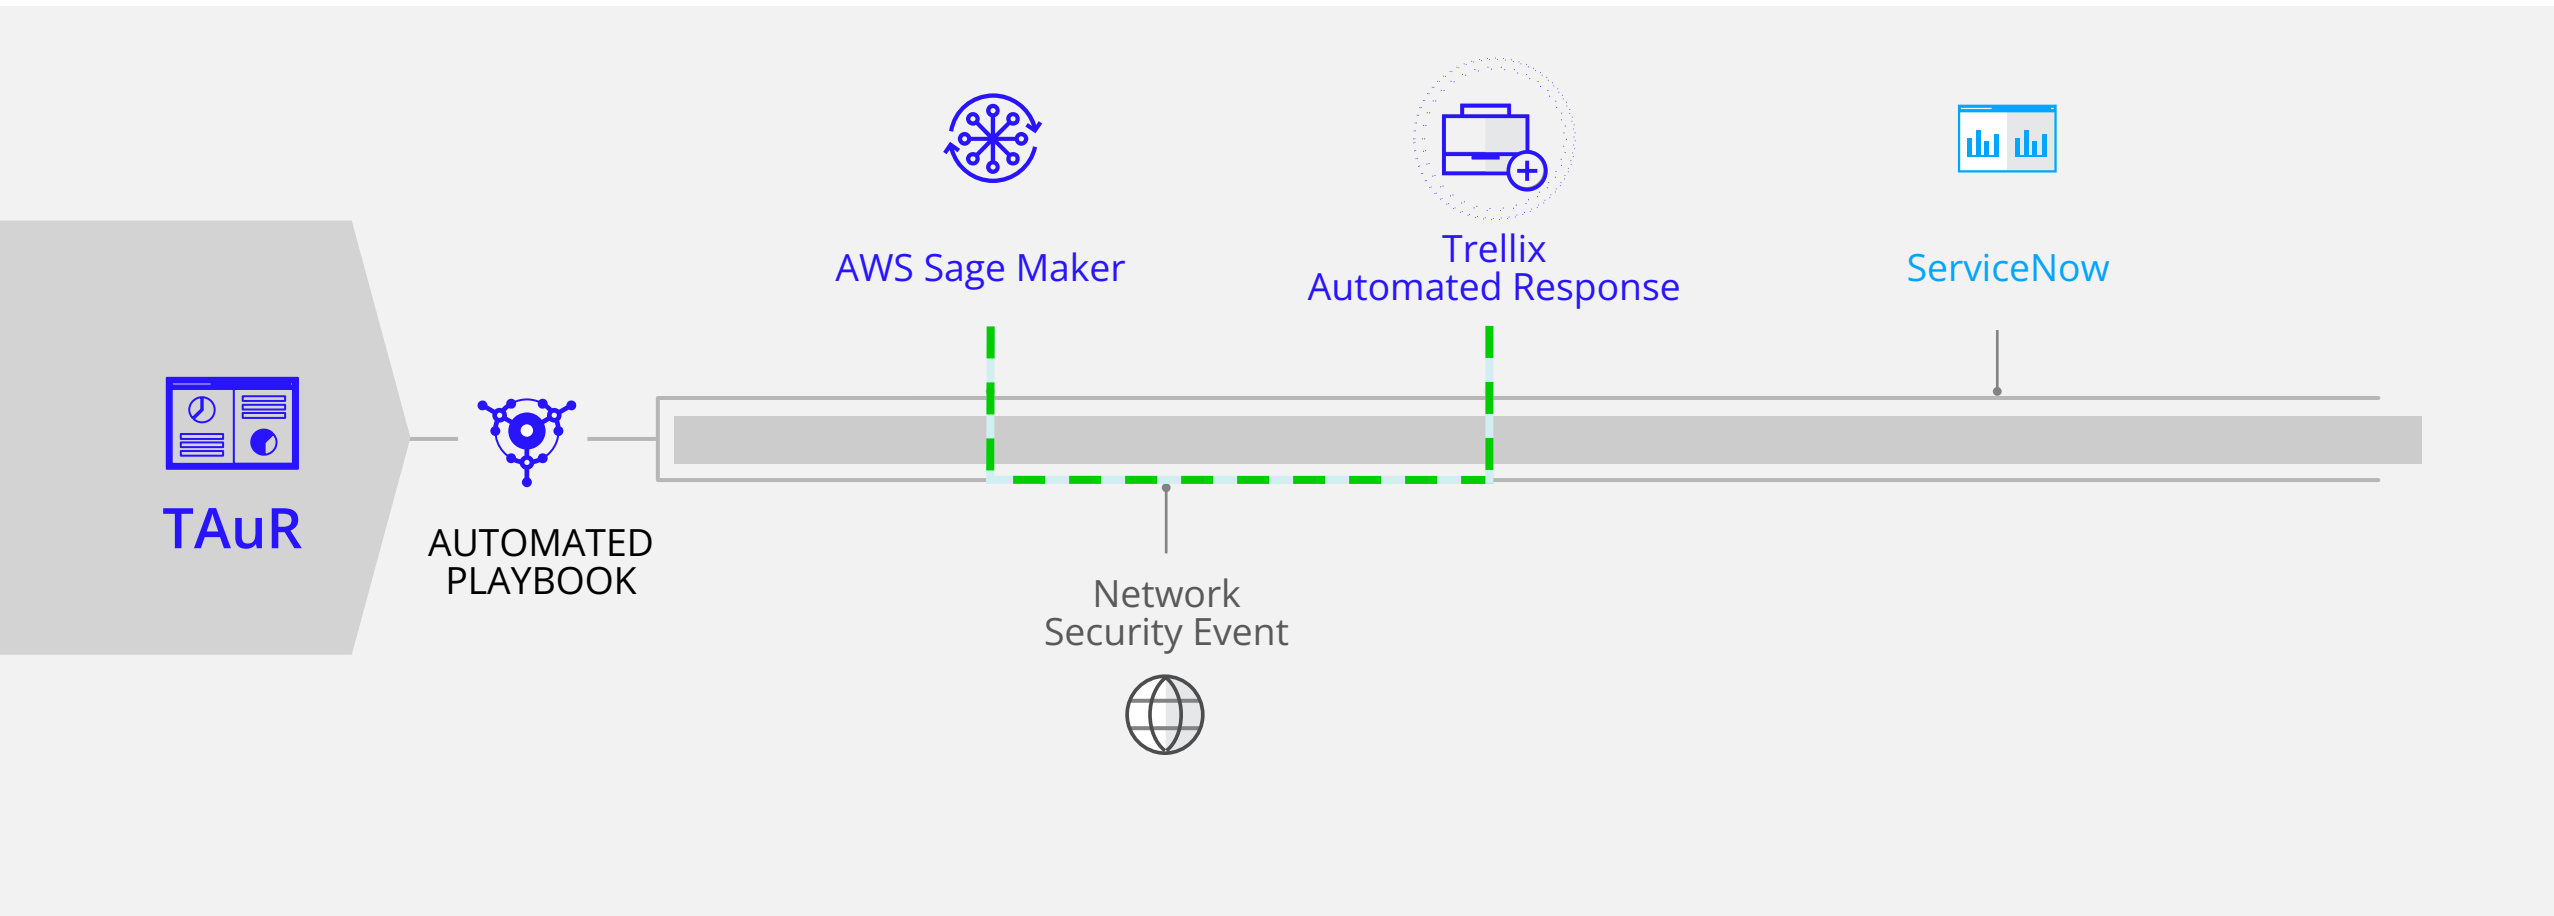
Playbook - Event Analysis with GenAI

Trellix Automated Response pulls that latest event from Trellix Network Security



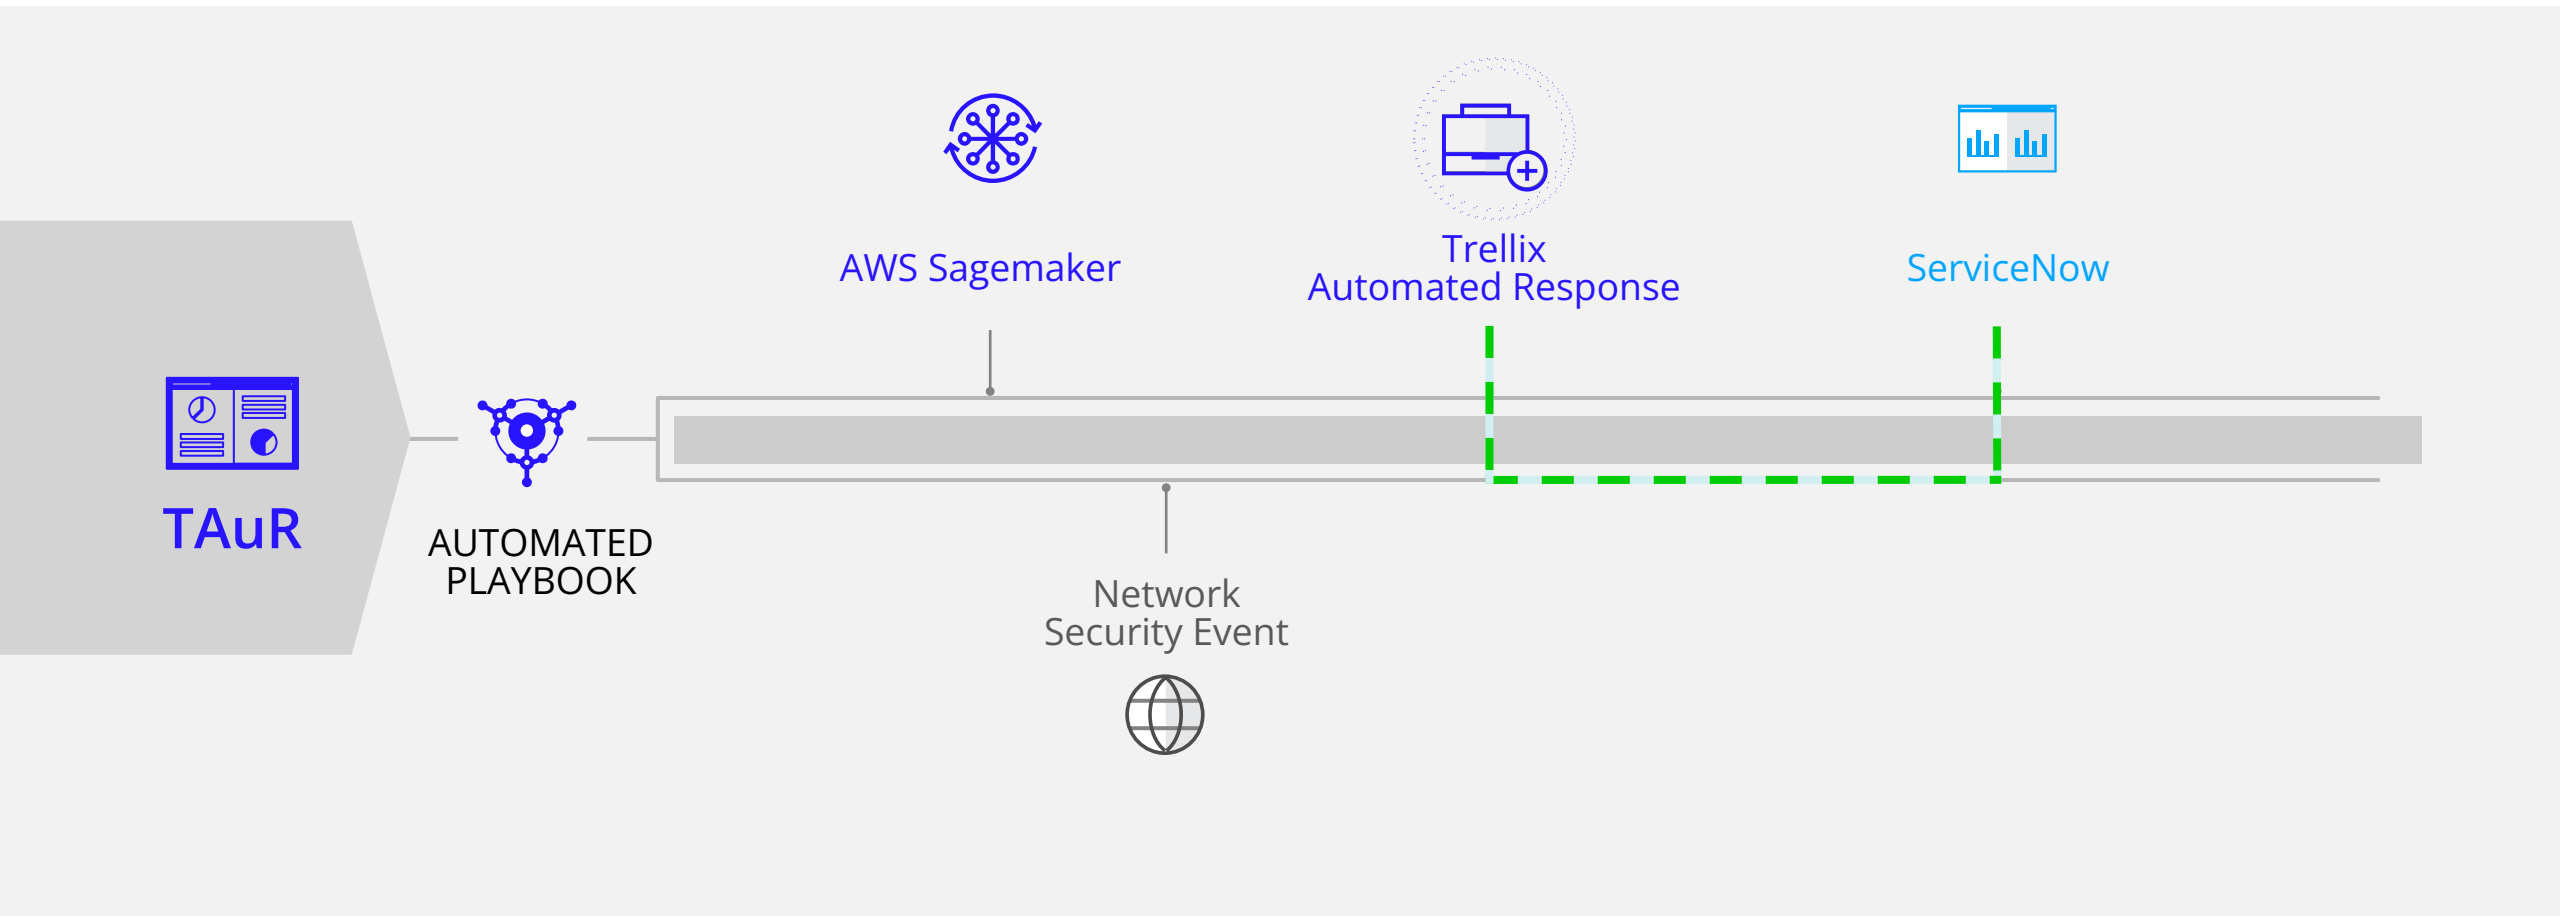
Playbook - Event Analysis with GenAI

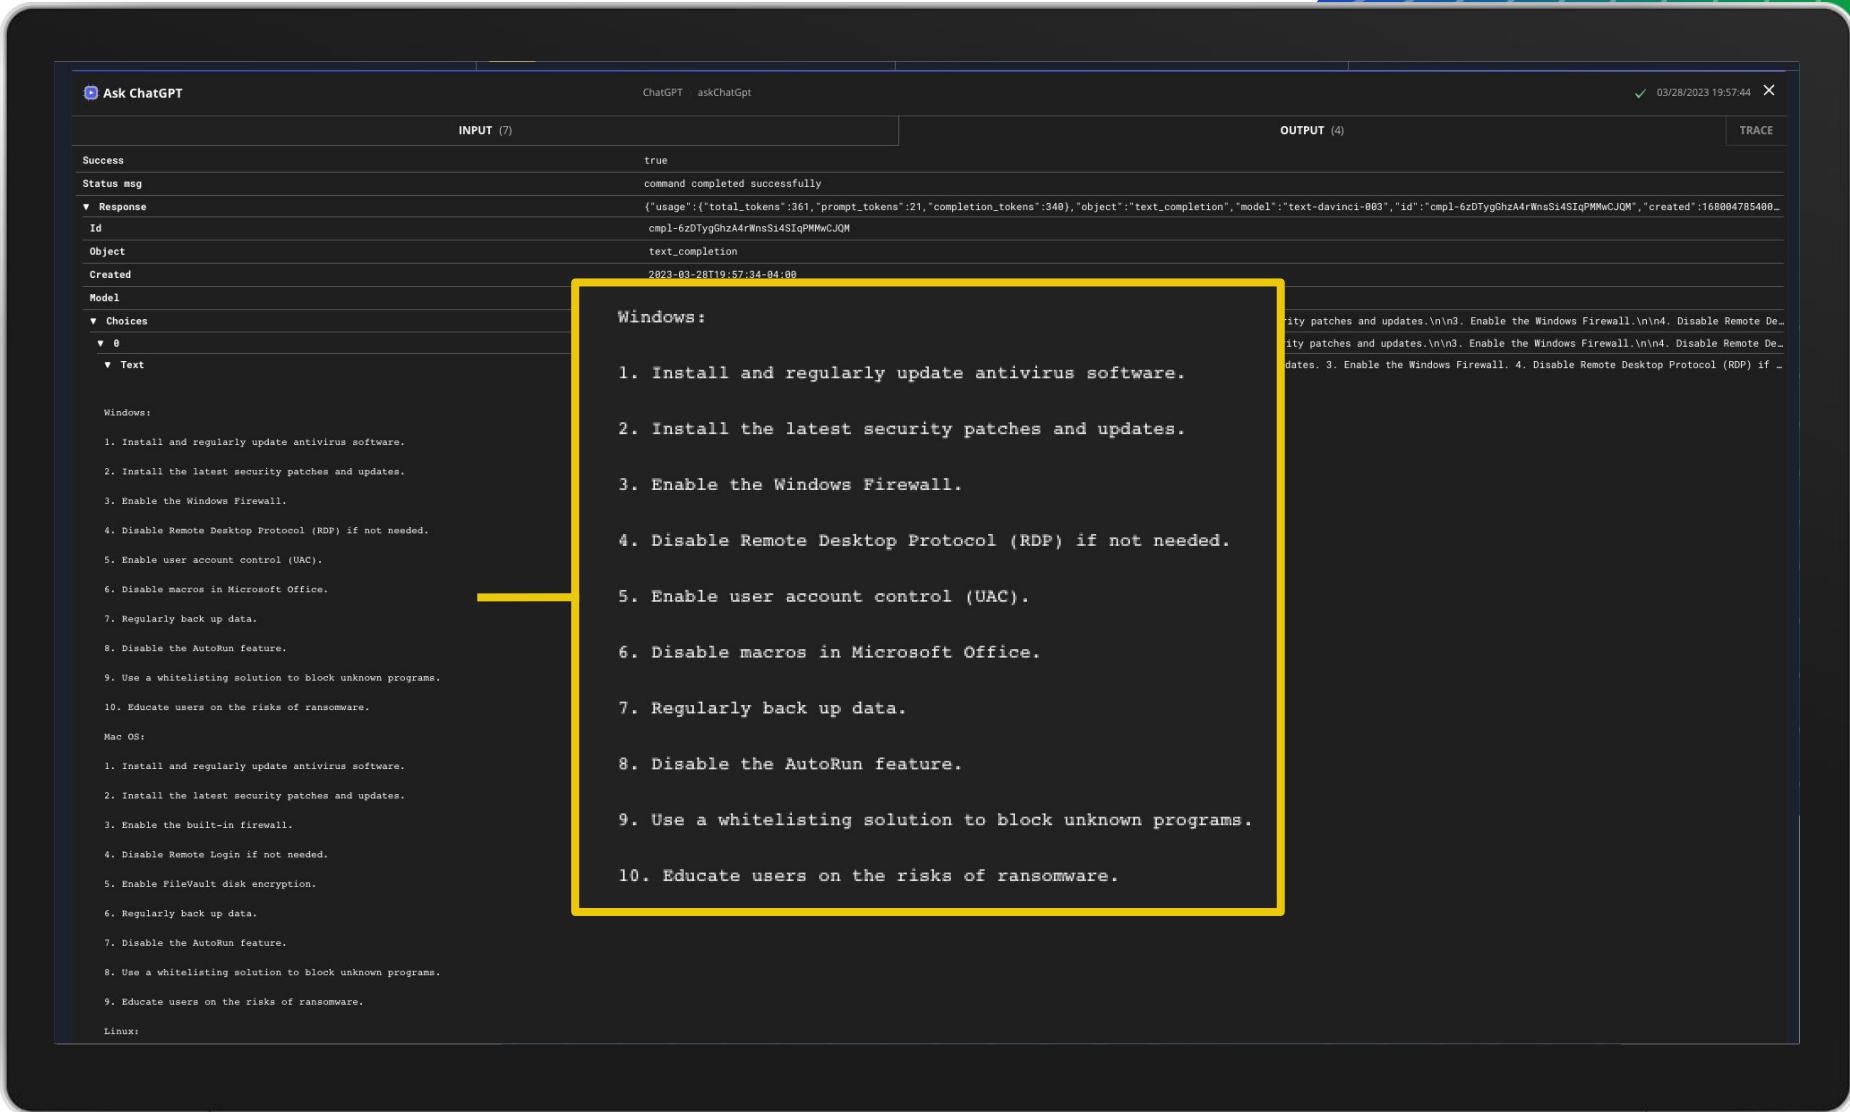
Trellix Automated Response initiates a playbook to call to AWS Sagemaker and pass over the event information to get a summarization, remediation steps, recommended protections



Playbook - Event Analysis with GenAI

Trellix Automated Response uses the results from AWS Sagemaker to create a ServiceNow ticket with the output and assign the ticket to the local network team







Trellix