RELIAQUEST®

Train Your Own AI Model for SecOps with **ReliaQuest GreyMatter**

How to leverage autonomous, self-learning agents for AI-driven security operations

Executive Summary

There's nothing more important in security operations than **speed**.

Containing threats within minutes—not hours or days—is the new standard. It's up to us in security to leverage all the latest tools to make this possible.

Agentic AI enhances security operations with faster, more accurate, and scalable threat detection, containment, investigation, and response while making autonomous decisions and recommendations.

Agentic AI differs from the more familiar Generative AI models (or large-language models [LLMs]), which might be considered passive AI systems. These models may simply execute instructions or perform narrowly defined tasks without independent initiative.

Agentic AI refers to autonomous artificial intelligence systems that have decision-making capabilities, functioning as independent "agents."

These AI systems can:

- Set goals
- Make decisions
- Take action

✓ Without constant human intervention.

Agentic AI is often designed to carry out specific tasks, optimize processes, or solve problems by analyzing data, learning from experience, and adapting to changing conditions.

To help organizations combat threats and move faster, ReliaQuest has built an AI Agent within its GreyMatter technology that empowers security teams to respond even more quickly and effectively without sacrificing visibility.

GreyMatter gives customers full control over training their own self-learning AI Agent, tailored specifically to their unique SecOps environment. By leveraging their own data, historical alerts, and real-time context, customers can fine-tune the AI to address their specific challenges, adapt to evolving threats, and optimize their workflows.

The GreyMatter AI Agent: Designed for SecOps

- **Built for Transparency**
- **Decades of Data and SecOps Expertise**
- **Access to Real-Time Data**
- **Privacy Built-In**

Autonomous Agent:

- Generates plans
- Collects additional insights
- Replans as needed

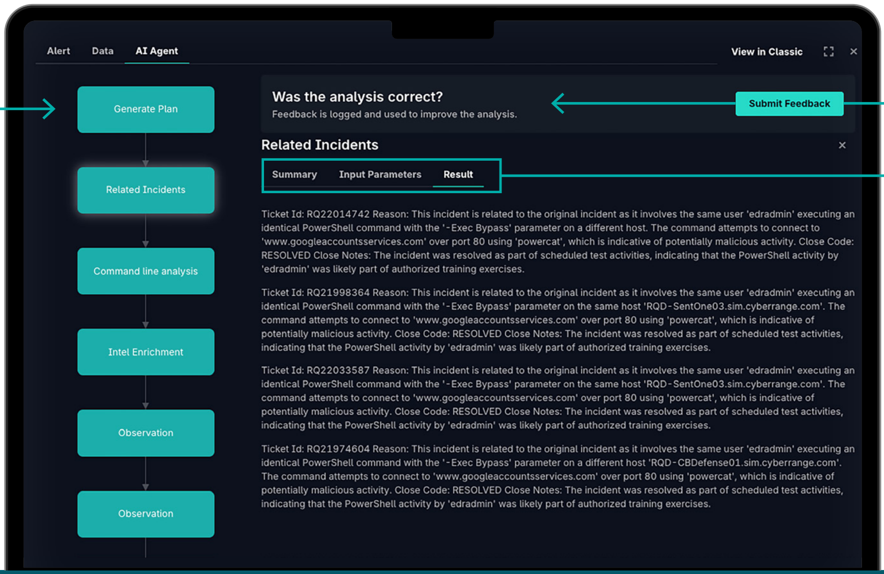

Was the analysis correct?

Feedback is logged and used to improve the analysis.

Related Incidents

Ticket Id: RQ22014742 Reason: This incident is related to the original incident as it involves the same user 'edradmin' executing an identical PowerShell command with the '-Exec Bypass' parameter on a different host. The command attempts to connect to 'www.googleaccountsservices.com' over port 80 using 'powercat', which is indicative of potentially malicious activity. Close Code: RESOLVED Close Notes: The incident was resolved as part of scheduled test activities, indicating that the PowerShell activity by 'edradmin' was likely part of authorized training exercises.

Ticket Id: RQ21998364 Reason: This incident is related to the original incident as it involves the same user 'edradmin' executing an identical PowerShell command with the '-Exec Bypass' parameter on the same host 'RQD-SentOne03.sim.cyberrange.com'. The command attempts to connect to 'www.googleaccountsservices.com' over port 80 using 'powercat', which is indicative of potentially malicious activity. Close Code: RESOLVED Close Notes: The incident was resolved as part of scheduled test activities, indicating that the PowerShell activity by 'edradmin' was likely part of authorized training exercises.

Ticket Id: RQ22033587 Reason: This incident is related to the original incident as it involves the same user 'edradmin' executing an identical PowerShell command with the '-Exec Bypass' parameter on the same host 'RQD-SentOne03.sim.cyberrange.com'. The command attempts to connect to 'www.googleaccountsservices.com' over port 80 using 'powercat', which is indicative of potentially malicious activity. Close Code: RESOLVED Close Notes: The incident was resolved as part of scheduled test activities, indicating that the PowerShell activity by 'edradmin' was likely part of authorized training exercises.

Ticket Id: RQ21974604 Reason: This incident is related to the original incident as it involves the same user 'edradmin' executing an identical PowerShell command with the '-Exec Bypass' parameter on a different host 'RQD-CBDefense01.sim.cyberrange.com'. The command attempts to connect to 'www.googleaccountsservices.com' over port 80 using 'powercat', which is indicative of potentially malicious activity. Close Code: RESOLVED Close Notes: The incident was resolved as part of scheduled test activities, indicating that the PowerShell activity by 'edradmin' was likely part of authorized training exercises.

Human-Validated Feedback

Transparent look into the mind of AI

Table of Contents

Executive Summary 1
Use Case: Agentic AI for Alert Investigations 1
Built by SecOps for SecOps 2
  Prompts Engineered by Security Analysts, Backed by Data Scientists 2
  Reinforcement Learning from Human Feedback 3
Years of Incident Response Data for Smarter AI 3
  Universal Translator to Normalize Data Across Security Tools 3
Unparalleled Transparency 4
Reduced Hallucinations 4
  Proprietary Knowlege Access 4
  Post-Generation Validation 4
Privacy Built-In 5
Conclusion 5
About ReliaQuest 5

Use Case: Agentic AI for Alert Investigations

This section highlights the step-by-step thought process of an AI agent designed to autonomously investigate security alerts. In this example, the AI agent generates a plan to investigate an incident involving Nmap execution. Nmap is a common tool used by threat actors to scan an organization's infrastructure to expose vulnerabilities they can exploit. The plan lays out a triage plan using AI tools available within GreyMatter that the executor can then act on. The example also showcases the output generated by the command line AI tool.

Incident Summary

This incident involving Rule ID 787 detects the execution of Nmap, a network reconnaissance tool, which was run by user "RELIAQUEST\\sa-chewell" on the host "AWUE-NP-DE01" to scan network ports on the range 172.28-29.108.62. Nmap is recognized as a utility commonly used for post-exploitation activities in cybersecurity incidents.

Triage Plan
(STEP 1)

Use the "Command line analysis" tool with the provided incidentId RQ20122203 to understand the context and purpose of the command line script that was executed.

Use the "Enriching artifacts" tool with the file hash f2496b4588c96a17aba7a6a9ac06dbaf69363dd8f2b4d008555aacc40867f78f to determine if the file associated with nmap.exe is known to be malicious or if it is a legitimate version of the tool.

Use the "Analyzing related incidents" tool with the incidentId RQ20122203 to find any historical incidents that are similar to the current alert, which could indicate a pattern of behavior or provide context for the usage of the tool.

Use the "IP Analysis" tool to examine the IP range 172.28-29.108.62.

Use the "Search technologies" tool with the incidentId RQ20122203 and a query to retrieve logs from the src_host AWUE-NP-DE01 for the time around the timestamp 2024-09-13T14:40:15.062Z to see if there are any other related events or commands executed by the user RELIAQUEST\sa-chewell.

Use the "Analyzing events" tool with the queryBatchUuid and index obtained from step 3 to further investigate the events, focusing on the command line activity, especially related to the output from the command line analysis tool, which analyzed the command flagged in the triggering artifacts.

If there are any other commands that look suspicious found from the "Analyzing events" tool, use the "Command line analysis" tool to analyse what the purpose and context of these commands were.

Use the "User lookup" tool to validate that the user's command line activity aligns with their job role.

Review the results from the previous steps to determine if the usage of nmap.exe was part of legitimate administrator activity or if it is indicative of malicious behavior. If the usage is in line with normal administrative tasks and there are no other suspicious indicators, it may be considered a false positive. If the usage is unusual or there are additional indicators of compromise, further investigation will be required.

Based on the findings, conclude whether the security incident is a false positive or a true positive and recommend any necessary response actions.

Command Line
(STEP 2)

The command is command line for nmap, which is a popular network scanning tool used for network discovery and security auditing. Here's a breakdown of the command and its options:

nmap: This is the command to run the nmap network scanner.

--reason: This option tells nmap to provide the reason why a specific port is set to a particular state (e.g., "reset" for closed ports).

--open: This tells nmap to only show open ports. Closed or filtered ports will not be displayed in the output.

-Pn: This option skips the discovery phase, where nmap tries to determine if the host is online before scanning it. It treats all hosts as if they are online, which is useful when scanning hosts that are blocking ping probes.

-p 22,88,389,636,464,443,3389,8000,8443,8051,8089,8501,139: This specifies the list of TCP ports to scan. The ports listed are commonly used for services such as SSH (22), Kerberos (88), LDAP (389), LDAPS (636), Kerberos change/set password (464), HTTPS (443), Remote Desktop Protocol (3389), and various other services that might be running on the additional ports listed.

172.28-29.108.62: This is the target specification for the scan. It indicates that nmap should scan the range of IP addresses from 172.28.108.62 to 172.29.108.62.

In summary, the command is instructing nmap to scan a range of IP addresses for open ports among a specified list of common service ports, without attempting to ping the hosts first, and to report the reason why each port is in a particular state. This type of scan is typically used to identify potential points of entry on a network and to understand the services that are exposed to the network, which is a critical step in a security assessment.

Built by SecOps for SecOps

At its core, the GreyMatter AI Agent is designed to meet the specific demands of security operations teams. Unlike generic AI solutions, GreyMatter is built with direct input from expert security analysts who understand the day-to-day challenges of defending an organization.

Prompts Engineered by Security Analysts, Backed by Data Scientists

Prompt engineering involves crafting specific instructions for AI agents to follow to ensure they produce the desired output.

Good prompting requires expertise and precision.

It is essential to craft advanced and detailed prompts, specific to the capabilities of the LLM, to guide the model towards the desired outcome.

For instance a simple prompt like "How many letters are in the word strawberry are 'r'?," now infamous online, yields an incorrect answer. However, a more detailed prompt that suggests ways to verify a response, such as **"How many letters in the word 'strawberry' are 'r'?** Verify your answer by checking each letter and then count the number of occurrences of each letter," can lead to a correct answer.

At ReliaQuest, prompt engineering is highly influenced by our expert security analysts and a comprehensive, well-documented cyber analysis methodology (CAM) that has been tuned and perfected using real-world data for over a decade.

ReliaQuest also incorporates proprietary prompting principles and techniques to ensure output from AI agents undergoes extensive verification to generate higher-fidelity outcomes. By employing advanced techniques such as Mixture of Prompts, Tree of Thought, Chain of Thought, and Graph of Thought, we've enhanced the quality of prompts and outputs. These techniques enable the AI Agents to handle complex tasks more effectively, ensure logical progression and coherence, and provide comprehensive and accurate responses.

Reinforcement Learning from Human Feedback

Incorporating human feedback is crucial for improving AI performance. ReliaQuest has a built-in feedback workflow for security analysts with domain knowledge to evaluate the output and steps taken by AI agents as a part of incident analysis.

This feedback loop allows customers to continuously train and refine their AI Agent, ensuring it evolves to handle emerging threats, adapt to new tools or workflows, and deliver increasingly accurate and actionable results. By combining ReliaQuest's expertise with customer-specific data, GreyMatter enables organizations to build an AI model custom to their security environment.

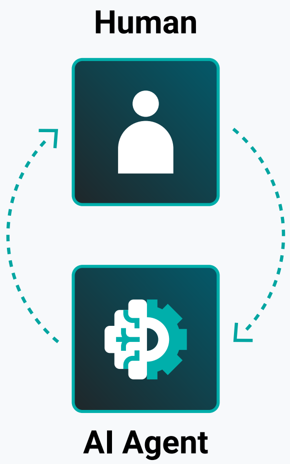

Years of Incident Response Data for Smarter AI

The effectiveness of an AI Agent depends on the quality of the data and methodologies it has access to. ReliaQuest pairs decades of historical alerts and incident response data with an extensive CAM. This CAM includes detailed information on various alert types, enrichment steps, and containment strategies, providing a robust foundation for customers to train their own AI Agents.

Using this foundation, customers can fine-tune their AI with real-time data from their own environments, creating an Agent that grows smarter and more effective over time. The result is a highly personalized AI solution that reflects the specific workflows, threat profiles, and operational realities of each organization.

Universal Translator to Normalize Data Across Security Tools

RELIAQUEST
GREYMATTER®

As part of the GreyMatter security operations platform, the Universal Translator normalizes real-time and historical data from multiple systems, including endpoints, SIEMs, and multiple clouds, enabling AI to process data and structure queries without depending on product-specific nuances.

For AI agents to be successful, system data must be normalized to a consistent ontology to allow for consistent, apples-to-apples comparisons. Security data from across all tools must be normalized to allow for scalable use for AI.

Because of this efficiency, ReliaQuest is well placed to scale its AI as it takes on new customers and their complex environments.

Unparalleled Transparency

To help customers develop trust in AI, ReliaQuest provides complete transparency into the way our AI Agent operates. Every step of a plan developed by the AI Agent—including data retrieved, inputs required, executed actions, failed steps, and observations—is displayed within GreyMatter as an internal monologue. This transparency allows users to interrogate and provide feedback on the AI's processes, fostering trust while facilitating improvements.

Reduced Hallucinations

One of the primary challenges with LLMs is their tendency to generate plausible sounding but incorrect information, a phenomenon known as hallucination.

Proprietary Knowledge Access

To combat this issue, ReliaQuest leverages Retrieval Augmented Generation (RAG) systems in real-time to provide LLMs with relevant information on historical incident response data, real-time customer context, and alert history to generate highly relevant outputs unique to each organization.

This is achieved with the following capabilities →

Integration with decades of proprietary cyber incident response data:

Our RAG system is tightly integrated with ReliaQuest's vast collection of cybersecurity data, including historical alert data, threat intelligence feeds, and customer-specific information, allowing AI models to seek additional information that may not be part of their training data.

Context-aware knowledge graph:

When processing an alert, the AI Agent also retrieves relevant customer-specific information from our data stores based on alert context. This might include similar past incidents from a specific customer's environment; known threat actor tactics; or specific customer environment details about assets, identities, and applications.

Augmented Generation:

The retrieved information is then used to augment the AI agent's knowledge, effectively constraining its responses to factual, relevant data. This process significantly reduces the likelihood of the model generating false or irrelevant information.

Post-Generation Validation

In addition to ensuring accuracy during the generation process, GreyMatter uses lightweight, automated quality assurance modules called "guardrails" to analyze the AI's outputs to confirm that they adhere to the specific requirements of the initial prompt.

For example, imagine a customer prompts the GreyMatter AI Agent to analyze an alert involving anomalous login activity and requests recommendations for mitigating the threat.

By integrating guardrails into the workflow, GreyMatter not only reduces hallucination but also ensures that the AI's outputs meet operational requirements with precision and reliability

The guardrails validate the output by ensuring:

The response includes specific recommendations aligned with the customer's environment (e.g., "quarantine the affected user account in Active Directory").

The response references relevant historical data from the knowledge graph (e.g., "this behavior aligns with a known tactic used by Threat Actor X in previous incidents").

No extraneous or hallucinated information is introduced, such as unrelated mitigation steps or fabricated data about threat actors.

Privacy Built-In

The GreyMatter AI Agent is built with data privacy and security in mind. Customer data is fully segmented, and each AI execution is context-specific, operating within the boundaries of the customer's data environment. In addition, by using privately hosted AI models and data-loss-prevention capabilities, we ensure that proprietary or sensitive data is not exposed to public AI models or to other customer environments.

Conclusion

ReliaQuest has harnessed decades of SecOps data to train generative AI and agentic AI models within its GreyMatter platform, making it uniquely suited for customers looking to augment their SecOps teams.

Pairing these AI capabilities with automation allows organizations to dramatically improve their security operations, reducing mean times to contain (MTTC) to 5 minutes or less while eliminating the burden of Tier 1 and Tier 2 security tasks.

> **What makes GreyMatter truly stand out is its ability to empower customers to train their own AI Agent, tailored specifically to their unique environments.**

About ReliaQuest

ReliaQuest exists to **Make Security Possible**.

Our Agentic AI-powered security operations platform, GreyMatter, allows security teams to detect threats at the source, contain, investigate, and respond in less than 5 minutes—eliminating Tier 1 and Tier 2 security operations work.

GreyMatter uses data-stitching, detection-at-source, AI, and automation to seamlessly connect telemetry from across cloud, multicloud, and on-premises technologies.

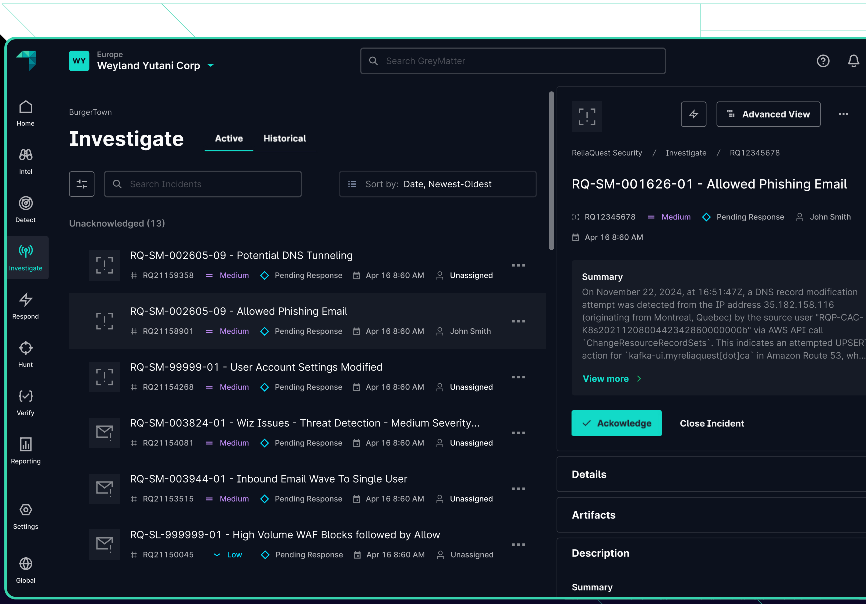

RQ-SM-002605-09 - Potential DNS Tunneling
# RQ21159358 Medium Pending Response Apr 16 8:60 AM Unassigned

RQ-SM-002605-09 - Allowed Phishing Email
# RQ21158901 Medium Pending Response Apr 16 8:60 AM John Smith

RQ-SM-99999-01 - User Account Settings Modified
# RQ21154268 Medium Pending Response Apr 16 8:60 AM Unassigned

RQ-SM-003824-01 - Wiz Issues - Threat Detection - Medium Severity...
# RQ21154081 Medium Pending Response Apr 16 8:60 AM Unassigned

RQ-SM-003944-01 - Inbound Email Wave To Single User
# RQ21153515 Medium Pending Response Apr 16 8:60 AM Unassigned

RQ-SL-999999-01 - High Volume WAF Blocks followed by Allow
# RQ21150045 Low Pending Response Apr 16 8:60 AM Unassigned

RQ-SM-001626-01 - Allowed Phishing Email
RQ12345678 Medium Pending Response John Smith
Apr 16 8:60 AM

Summary
On November 22, 2024, at 16:51:47Z, a DNS record modification attempt was detected from the IP address 35.182.158.116 (originating from Montreal, Quebec) by the source user "RQP-CAC-K8s20211208004423428600000000b" via AWS API call `ChangeResourceRecordSets`. This indicates an attempted UPSER action for `kafka-ui.myreliaquest[dot]ca` in Amazon Route 53, wh...
View more

ReliaQuest is the only cybersecurity technology company that delivers outcomes specific to each organization's unique architecture, technology and business needs.

With over 1,000 customers and 1,200 teammates across six global operating centers, ReliaQuest Makes Security Possible for the most trusted enterprise brands in the world. **Learn more at www.reliaquest.com**.

RELIAQUEST®

reliaquest.com
800.925.2159
info@reliaquest.com

Copyright © 2024 ReliaQuest, LLC. All Rights Reserved. ReliaQuest, GreyMatter, Digital Shadows, and all related logos, product and services names, designs, and slogans are trademarks or registered trademarks of ReliaQuest, LLC or its affiliates or licensors. All other product names or slogans mentioned in this document may be trademarks or registered trademarks of their respective owners or companies. The ReliaQuest software, platform, portal and its entire contents, features, and functionality are owned by ReliaQuest, LLC and its affiliates. These materials are protected by United States and international copyright, trademark, patent, trade secret, and other intellectual property or proprietary rights laws. All other information presented is provided for informational purposes with no representations or warranties provided of any kind and should not be relied upon for any purpose. ReliaQuest has no obligation to amend, modify, or update the information contained in this document in the event that such information changes or subsequently becomes inaccurate. Printed in the USA.