TAG

USING ARTIFICIAL INTELLIGENCE FOR SECOPS AUTOMATION

DR. EDWARD AMOROSO,
CHIEF EXECUTIVE OFFICER, TAG[1]

SWIMLANE

USING ARTIFICIAL INTELLIGENCE FOR SECOPS AUTOMATION

DR. EDWARD AMOROSO, CHIEF EXECUTIVE OFFICER, TAG[1]

This TAG analyst report explains how modern SecOps automation methods can leverage artificial intelligence (AI) technology to improve effectiveness. The AI-enhanced security automation platform from Swimlane[2] is used to demonstrate the concepts in the context of a commercial-available solution.

INTRODUCTION

Automation of the security operations center (SOC) is a high priority for enterprise teams who must now deal with increasingly automated and AI-driven offenses. Threat actors are now using AI to exploit vulnerabilities, automatically run through attack paths, craft phishing emails, impersonate speech, and engage in misleading chat conversations. These methods make offensive cyber threats much more difficult to prevent, detect, or mitigate.

The traditional approach of following manual processes in the SOC to collect data, make human-time decisions, and perform manual response handling must now be augmented with an automated means for keeping up with the required pace of security. As one might expect, modern AI solutions using machine and deep learning offer great promise toward driving greater levels of automation for cyber defense.

The good news is that SOC teams have understood the need for automation for some time, and most have engaged the partnership of a commercial vendor to help reduce manual work. For many of these practitioner teams, first-generation Security Orchestration, Automation, and Response (SOAR) platforms have been helpful, but the time has arrived for SOAR platforms to evolve toward an even higher level of scalability, effectiveness and simplicity.

In this report, we show how SOAR platforms have evolved to next-generation AI-enhanced security automation solutions designed to deal with the threat challenges and intelligent automation opportunities present in the modern SOC. We demonstrate these concepts by showing how commercial vendor Swimlane is implementing generative AI-based capabilities to enhance automation for the SOC as well as across the enterprise security ecosystem.

BRIEF OVERVIEW OF SWIMLANE AND THE SOAR MARKET

Swimlane was founded in 2014 to focus on the growing need for SOC teams to have the scalability and flexibility needed to keep pace with the velocity of alerts coming from their security information event management (SIEM) platform and related detection tools. By 2017, this prompted the creation of the SOAR market category referenced above and the first-generation of commercial SOAR tools.

Swimlane's early years were marked by rapid innovation and development, driven by their commitment to creating a scalable, flexible, and fast automation engine for SOC teams to integrate and automate anything. The goal of the *Swimlane Turbine* platform has always been to alleviate the burden of repetitive tasks on security analysts, allowing them to concentrate on more strategic activities.

Swimlane Turbine gained traction with SOC teams for its ability to integrate with any application programming interface (API), thus creating improved visibility and data flow across an enterprise security infrastructure. This integration capability became a cornerstone of Swimlane's success, enabling organizations to maximize the value of their existing security investments by extending automation use cases beyond that of traditional SOAR tools.

As Swimlane expanded, the platform evolved with a cloud-native and low-code architecture, as well as robust analytics and reporting capabilities. These enhancements allowed SOC teams to identify patterns, understand threat landscapes, and make more informed decisions. The platform's user-friendly interfaces and customizable workflows further solidified its appeal, ensuring that it could be tailored to meet the specific needs of diverse organizations.

Swimlane's focus on case management and incident response has also been a defining feature. The platform facilitates the efficient handling of security incidents from detection through resolution, ensuring that no step is overlooked. This systematic approach not only improves response times but also aids in cross-functional collaboration, documenting and learning from past incidents, thereby enhancing future preparedness.

It should thus come as no surprise that Swimlane would leverage modern AI technology to continue this journey toward support for *more trusted, secure, and transparent use of AI in the SOC*. In the sections below, we explain and illustrate this effective use of AI to transition Swimlane customers from existing SOAR technology toward a more advanced and effective AI-enhanced security automation approach.

HOW AI-ENHANCED SECURITY AUTOMATION WORKS

To understand how AI is being integrated, it helps to review additional features in the existing Swimlane Turbine solution. First, the platform can handle vast data sources, which any observer of AI technology will understand to be a prerequisite for implementing machine or deep learning. Data is processed efficiently using a novel multi-agent architecture that can serve as a foundational basis for the evolving prompt engineering approaches used in AI.

The emerging Swimlane Hero AI is centered around main personas – namely, the SOC analyst, the orchestrator, and the manager. Hero AI is designed to empower the SOC analyst to increase their effectiveness (perhaps taking Tier 1 analysts to Tier 3 capability). This is done by simplifying how SOC Analysts communicate with the data (essential in AI) and building a set of actionable and relevant suggestions.

Swimlane Hero AI exists to empower the orchestrator role by increasing their operation's effectiveness via reducing the time needed to create automations and reducing the cognitive complexity of their tools. Practitioners using AI tools in the SOC will recognize this task of defining automations in workflow, and any effort to reduce the time and complexity of this task should be well appreciated.

Ultimately, Swimlane Hero AI was created to partner with and empower the manager role by increasing visibility and improving their ability to support planning. Managers feel considerable pressure these days to demonstrate that AI can be integrated into operations. As such, it is important to ensure that their ability to pull management data and generate reports is supported in any deployed AI-based platform.

Throughout its support for its SOC team partners, Swimlane Hero AI knows that it must ensure the full privacy of customer data. This includes customer data never leaving the Turbine Cloud, never being shared with other customers, and never used for model training. These are important considerations for any enterprise team concerned with establishing guard rails for their AI usage. With Swimlane Hero AI, the guard rails are built in (more on this below).

FROM SOAR TO AI-ENHANCED SECURITY AUTOMATION

Practitioners understand that SOAR has always been focused on SOC automation use cases. These tools have primarily been adopted by mature SOC teams who have the developer resources needed to build and maintain integrations and playbooks (e.g., Python-based). Low-code innovations have also made cybersecurity tools like SOAR more approachable, and the addition of AI is poised to revolutionize the industry once more.

As a result, the integration of AI into the Swimlane Turbine platform should come as no surprise to existing or prospective customers. In fact, a range of common SOC automation use cases are already being enhanced using Swimlane Hero AI. The benefits of AI will further extend into adjacent use-cases including the securing of operational technology (OT) systems and reducing business fraud.

The term *AI-enhanced SecOps* has emerged in the industry to describe integration of AI into the SOC infrastructure, and this designation accurately describes the Swimlane approach. SOC practitioners should consider adopting such technologies to not only support existing SOC tasks, but also to create new opportunities for analysis and response that might not have been possible without the power of AI models to process large volumes of data.

FURTHER UNDERSTANDING SWIMLANE'S USE OF AI

To further understand how Swimlane is leveraging AI to improve SOC automation, it helps to go through the various functions supported in the commercial platform. As would be expected, this includes extensive use of AI models for processing, as well as the deployment of large language models (LLMs) to enhance the means by which SOC analysts can query data to obtain actionable guidance on prevention, detection, and response tasks.

Hero AI Support for Data Privacy and Security Standards
As suggested above, Hero AI prioritizes data privacy and security by leveraging privately hosted models, in particular, the Swimlane LLM. Swimlane ensures that no sensitive customer data is stored centrally. Only metadata related to model usage and performance is stored centrally, and customer data is never used for training of privately hosted shared models nor is it stored in any shared log files. Instead, data processed by AI/ML models in Swimlane is securely stored in a dedicated database instance for each customer, ensuring complete data segregation.

This data remains the property of the client, who may grant permission for Swimlane to use it solely for model validation. Furthermore, any trainable models that utilize customer data are enabled, trained, and served only with explicit customer consent and are exclusively used for that customer, never shared with others. These data privacy measures reinforce Swimlane's commitment to maintaining the confidentiality and integrity of customer information.

Swimlane AI-Based Data Processing

Swimlane Hero AI support begins with its handling of data. Any reader working in the area of AI will understand the essential role that high-volume data plays in establishing accurate models. Swimlane thus collects data in the SOC from virtually every possible source both internally and externally. The AI benefits from such coverage including from standards bodies such as MITRE and NIST, as well as other security knowledge bases.

Internal sources of data are also important to AI-based processing and Swimlane Turbine ingests from a wide variety of systems. This includes structured data from around the enterprise as well as more unstructured data sources such as case notes, alert details, data from knowledge base articles, metadata about playbooks, and many other sources. The integration of such unstructured provides valuable context for the Swimlane Hero AI algorithms.

Hero AI's machine learning algorithms were selected to optimize handling of these various data sets for training, as well as to handle the use-cases that the company has been supporting for many years with SOC customers. Buyers of AI-based tools should always be certain to understand how the selected platform ingests, utilizes, and leverages data since this will determine the overall effectiveness of the AI approach.

Swimlane Learning Algorithms

Swimlane employs state-of-the-art AI technology across its Turbine platform. For example, algorithms known as AI classification support the security process. These algorithms include both binary and multi-classification methods. In multiclass classification, each record belongs to one of three or more classes, and the algorithm's goal is to construct a function which, given a new data point, will correctly identify the class into which the new data point falls.

In addition, Swimlane utilizes methods known as LORA (Low-Rank Adaptation) and qLORA (Quantized Low-Rank Adaptation) to fine-tune Hero's large language models (LLMs) as part of so-called Parameter Efficient Fine Tuning (PEFT), which is designed to train the SOC security models to operate more efficiently than traditional methods. This is key to detection of new zero-day issues that might not be included in training data.

The use of LLMs in the Swimlane platform is to create a processing environment in which the automation not only complements the decision-making of the human SOC analysts, but that also creates a model that is more knowledgeable about the overall security domain. This includes the ability to tune processing based on the analyst's specific domain, as well as the ability to predict the next steps that are likely to be taken by an adversary.

Transformer-based AI models are also employed in the Swimlane platform. These are types of neural network architectures that transform an input sequence into a proper output. This is accomplished in the Swimlane solution through learning context and by tracking relationships between input sequence components. The result is that embeddings of text input to the LLM can be used to generate accurate security-related guidance for the SOC analyst.[3]

ADDITIONAL DETAIL: RAGS AND MULTI-AGENT LLMS

Rapid developments and advances in Large Language Models (LLMs) have spurred considerable recent research by expert teams, including at Swimlane, on the best ways to utilize AI for answering user questions. Most readers will already know that the most basic way for a human or other entity to communicate with an LLM is to ask it a question directly, or to send it along with certain instructions for the expected response.

A challenge does arise, however, if a user sends questions to the LLM that presume access to data or information accessible from some internal or non-public knowledge base. This challenge has spurred a new technique known as *Retrieval Augmented Generation (RAG)*. The RAG method involves adding relevant context to the prompt which can significantly increase the quality of the response.

Retrieval of the context using RAG for a question sent to an LLM usually works as follows: The internal knowledge base is first split into smaller pieces called chunks. The chunks are indexed by a separate model, called an *embeddings model*, with a numeric multi-dimensional vector, representing its semantic meaning. When a user asks a question, the question is first sent to that embeddings model to obtain its numeric vector representation.

The vector representation of the user question is then compared with numeric representation of the chunks, and the chunks with the closest vector representation are being selected. The original user question is sent to the LLM along with the content of the closest chunks from an internal knowledge base. This approach improves on direct responses coming from an LLM, although even state of the art technology might struggle to answer more complex questions.

It is worth mentioning that an even more advanced way of answering user questions is being examined at Swimlane. The method, which is known as a multi-agent LLM, involves building a system of several communicating AI agents. This could be a system of two or more agents, where one agent performs orchestration, focusing on building a plan for answering user questions, and other agents, supporting various types of RAG answer questions to the internal and external knowledge bases.

Operational Use of the Swimlane Function

Let's review how these advanced AI functions and LLM support in the Swimlane Turbine platform can produce use-cases that illustrate the effectiveness of combining the human experience and ingenuity of the SOC analyst with the power of the automated platform. A good starting point is to review the types of inputs that are likely to be provided to the Swimlane LLM by the SOC analysts during an incident or case investigation.

It is possible that the analyst might present to the Swimlane Turbine platform a series of artifacts related to the situation at hand. These might include both relevant and potentially non-relevant information such as articles, links, analyst questions, playbooks, and data about applications, networks, and systems that might be engaged in the incident. SOC analysts will recognize how this process can be tough to navigate manually and how it can be so time consuming to handle.

The result of the Swimlane Hero AI processing is that a series of recommendations will be provided based on the well-known means by which machine and deep learning algorithms (as outlined above) can ingest, analyze, and process data to offer actionable and relevant guidance. As one would expect, the human analyst is key to interpreting Hero AI's guidance and to determine the best means for taking appropriate action.

ECONOMIC BENEFITS OF AI-ENHANCED SECURITY AUTOMATION

It is worth spending some time reviewing the economic benefits that come with such capabilities, since it should be clear that the AI-enhanced security automation removes friction from manual, time-consuming, and hence expensive processes. Reviewing at a high level the types of expense savings that will accrue for the typical SOC reveals various estimates of how Swimlane AI-enhanced security automation will improve an organization's financials.

If we take the perspective that expense costs in a SOC can be expressed in units of productivity, especially as they relate to the need to augment employee staff with consultants and contractors, then by improving productivity by 20% utilizing AI and automation, then it is certainly possible that for every five staff in the SOC, one could be either redeployed to another task, or a contractor could be removed.

Extrapolating the math, if a SOC had twenty staff working on a regular basis, consisting of the usual mix of tiered analysts, employee staff, consultants, and contractors – and if we presume the average salary for one SOC staff to be USD $250K,[4] then the staff budget for that SOC would be USD $5M per annum. Reduction of 20% would save a million dollars which would easily cover any new license expenses for the AI-enhanced security automation platform.[5]

ACTION PLAN

The action plan recommended for SOC teams and their management is that the time has come to be serious about the full use of dedicated automation tools to optimize efficiency and effectiveness. And the use of AI ensures that this deployment, and the requisite investment required, will maximize their return on investment by reducing the complexity of legacy automation tools, and extending the benefits of AI and automation into new use-cases beyond the SOC.

As should be obvious, our recommendation here is that Swimlane offers a high level of transparency, security, and trust in their platform and would be a sensible choice to include in all source selection in this area. As always, TAG analysts are available to assist Research as a Service (RaaS) customers in the review, assessment, and decision-making process for all aspects of cybersecurity and artificial intelligence.

---

[1] TAG Infosphere provides research and advisory in cybersecurity, artificial intelligence, and climate science/sustainability for enterprise teams, government agencies, public policy lawmakers, academic researchers, and commercial vendors. See https://www.tag-infosphere.com/.

[2] Swimlane is a cybersecurity company that empowers security teams through automation. It is the largest and fastest-growing pure-play security automation company. See https://swimlane.com/ for more information.

[3] The public Swimlane website is an excellent source of additional information on how Swimlane employs AI-based technology to enhanced SOC productivity and effectiveness. See https://swimlane.com/.

[4] Obviously, it is not feasible to provide an accurate general estimate for SOC staff that covers all sectors, sizes of companies, and geographies. If anything, USD $250K would be low in most SOC environments, but readers can fill in their local estimated values.

[5] Readers should work a local calculation to determine a suitable return on investment (ROI) for the Swimlane automation. Replacement of legacy SOAR platforms could be worked into the calculation as would the many other qualitative ROI elements that would improve the day-to-day work activities of SOC analysts.

ABOUT TAG

TAG is a trusted research and advisory company that provides insights and recommendations in cybersecurity, artificial intelligence, and climate science to thousands of commercial solution providers and Fortune 500 enterprises. Founded in 2016 and headquartered in New York City, TAG bucks the trend of pay-for-play research by offering unbiased and in-depth guidance, market analysis, project consulting, and personalized content—all from a practitioner perspective.

IMPORTANT INFORMATION ABOUT THIS DOCUMENT

Contributors: Edward Amoroso

Publisher: TAG Infosphere Inc., 45 Broadway, Suite 1250, New York, NY 10006.

Inquiries: Please contact Lester Goodman at lgoodman@tag-cyber.com to discuss this report. You will receive a prompt response.