SentinelOne

Four Questions CISOs are Asking About Generative AI

eBook

Table of Contents

One Step Closer to Autonomous Security Operations

Generative AI is a Transformative Step in the Evolution of Human-Machine Interaction in Cybersecurity

AI already has a long history of use in security and is deeply embedded in SecOps workflows. Supervised machine learning is commonly used to scan scripts to detect malware while unsupervised machine learning is used to correlate signals, and detect suspicious activity. Historically, AI's impact on SecOps has been focused on detection. Generative AI is a type of artificial intelligence that can create new content such as text, images, videos, or audio from a prompt based on a large wealth of training data. This new technology brings exceptional capabilities in terms of **natural language understanding** and **code generation** and offers an opportunity for transformation that spans the breadth of the SecOps workflow. Our focus in this piece is on **large language models (LLMs)**, a subset of generative AI focused on producing human-like text, and their impact on security operations. We believe LLMs can help give time back to analysts by **automating mundane tasks** and augmenting the analyst's capacity to **search, correlate, query, and contextualize**.

Many CISOs are investigating this technology and seeking to grasp the implications – both opportunities and hazards. Will the risks of AI in the hands of adversaries outweigh benefits of AI automation in the hands of security teams?

Our view is cautiously optimistic. We believe generative AI tips the balance in favor of defenders by giving time back to overstretched security teams. That said, it's important to be strategic about adopting use cases. The greatest opportunities lie at the intersection of generative AI's strengths and the challenges facing SOCs today:

63%
of security practitioners are experiencing some level of burnout[1]

71%
said their organization had been impacted by the cybersecurity skills shortage (up from 57% in 2021)[2]

56%
of large companies handle 1,000+ security alerts each day[3]

63%
of organizations consider duplicate alerts a moderate to significant challenge, while 60% view false positives similarly[4]

Generative AI's most transformative impact on **human-machine interaction** is the democratization of cybersecurity, where the barrier to entry for complex tasks is lowered. This is made possible through generative AI's natural language interface, capable of taking a prompt, executing complex queries, and serving results in an intuitive format. This uplevels junior analysts and enables experienced analysts to focus on higher-impact, strategic tasks.

[1] Voice of the SOC, Tines, 2023

[2] The Life and Times of Cybersecurity Professionals Volume VI, ESG, 2023

[3] 56% of Large Companies Handle 1,000+ Security Alerts Each Day, Dark Reading, 2020

[4] CSA Official Press Release, CSA, 2024

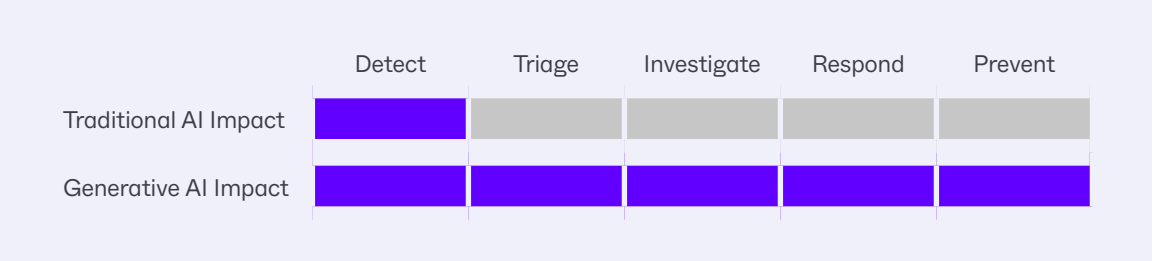

In this eBook, we will offer a perspective on the implications of generative AI in cybersecurity, both the immediate opportunities and long-term transformations. We'll do this by answering the four questions we are hearing most from CISOs today:

1 | What Impact is Generative AI Having on the Threat Landscape?

2 | What are the Top Use Cases for Generative AI in Cybersecurity?

3 | How can Organizations Evaluate Generative AI Offerings and Vendors?

4 | What are the Implications of Generative AI for the Future Role of the SOC?

1 | What Impact is Generative AI Having on the Threat Landscape?

A Weapon for Good or a Fearful Asymmetry?

LLM creators have gone to great lengths to build protections into their models; indeed, generative AI appears almost unique among technological innovations in terms of the depth of thinking and investment directed towards considering the social implications and implementing guardrails to promote safety and reduce misuse. Despite these laudable efforts, generative AI does introduce new risks. These tools are already being used by **cyber-attackers as well as defenders**. The most harmful effects may not necessarily fall under the category of "jailbreak and misuse." Adversaries can benefit from these tools in the same ways as ordinary users – helping them code faster, find relevant information, and accelerate mundane tasks.

So far, attackers have mainly use generative AI to dramatically increase the quality and sophistication of **social engineering**. Generative AI tools can produce convincing phishing messages at scale, in any language, even if the hackers themselves lack fluency. This means the "red flags" which have been the focus of much social engineering training are now less obvious. The rise of "**deepfakes**" has also enabled adversaries to convincingly imitate the voice or even video appearance of senior executives from a small sample, which can then be used to target humans or spoof biometric authentication systems. These capabilities, alongside phenomena like Ransomware-as-a-Service (RaaS), have **lowered the barrier to entry** for malicious activity. The consequence of this is a splintering of the threat landscape, with a proliferation of mid-level players frustrating efforts to maintain quality intelligence on threats.

There have been some tentative moves by malicious actors towards using generative AI for more technical use cases beyond social engineering.

Examples include:

- **LLM-Enhanced Scripting Techniques** LLMs unlock huge time savings for coding tasks of all types. While most models have protections against writing malicious script, it is nonetheless true that LLMs can be useful to adversaries to help write code more quickly, and meet complex parameters.
- **LLM-Enhanced Anomaly Detection Evasion** LLMs have been used to develop code that evades common detection mechanisms.
- **LLM-Assisted Vulnerability Research** Powerful and publicly available semantic search tools have been exploited by adversaries to gather information on publicly-known vulnerabilities. These tools also have broader reconnaissance uses, such as helping organizations identify and research potential targets.

AI in the Headlines

Fraudsters Cloned Company Director's Voice in $35 Million Heist[1]

Hackers Use AI to Bypass Bitfinex's Biometric Authentication System by Injecting Fake Video Stream[2]

Generative AI is Social Engineering on Steroids[3]

[1] Forbes, 2023

[2] Deloitte, 2023

[3] Inc, 2023

Our overall observation is that adversary use of generative AI is still nascent outside the realm of social engineering. However, this could change in the future. We see several theoretical capabilities which could evolve into meaningful threats:

- **AI-powered polymorphic malware** is a type of malicious software that leverages artificial intelligence to constantly change its code, making it difficult to detect and defend against.
- **Generative AI agent-based techniques** would use AI to string together several complex automated actions.
- **Accelerated exploit discovery** would use generative AI to creatively and adaptively penetration test an environment for weaknesses.

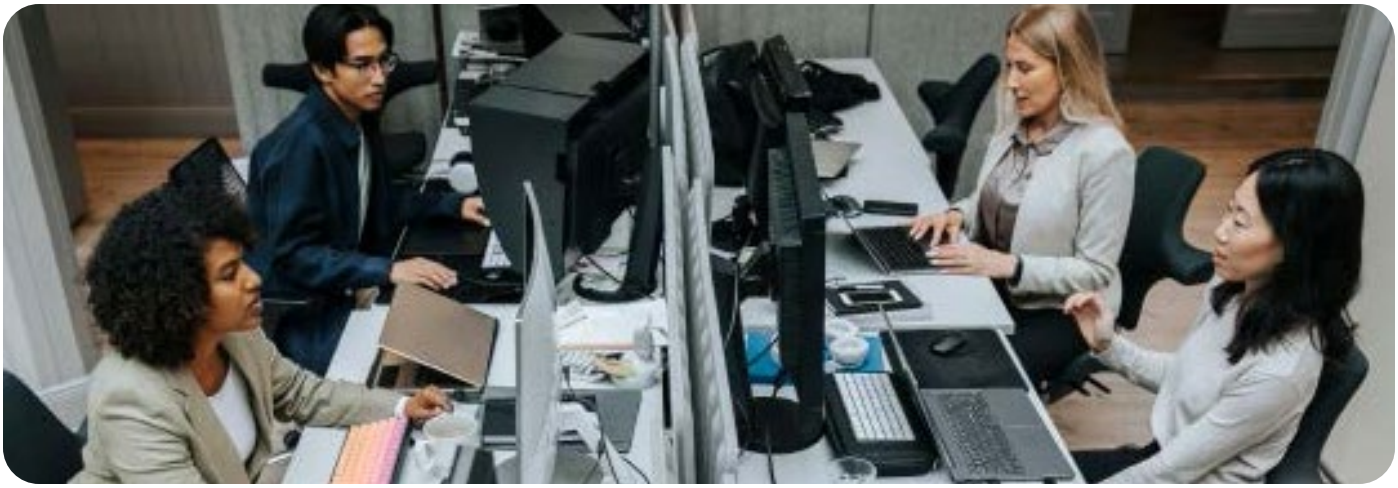

A New Attack Surface

Many organizations adopted generative AI solutions very quickly, as the technology gained considerable media attention. At the same time, CISOs have been working to secure this growing use of generative AI, asking themselves questions like:

- Is adequate data protection in place, as these models are exposed to large quantities of proprietary data?
- Is sensitive or proprietary data being loaded into LLMs in a way that it could be leaked or that I could lose control of it?
- Is the model protected against misuse or jailbreak?
- How can I help prevent generative AI-based supply chain threats in the software development lifecycle?
- Should I be applying zero trust principles to AI tools, applications, and platforms, or enable continuous monitoring and dynamic access controls?

Generative AI can exacerbate a risk already plaguing many organizations – **overprivileged access**. Most organizations today still don't have the level of granular control over their data that they would like. This is partly because organizations need to share and collaborate, often with third-party users, and can have thousands of individual users sharing files with peers or placing files in shared folders that are often accessible to large groups. In this context, it is difficult to prevent things from being overlooked and ensure ordinary users maintain good hygiene. Generative AI makes this problem worse because many organizations now have highly sophisticated **semantic search tools**. Where previously a user might stumble on something they shouldn't, now they have a powerful search tool that can expose overprivileged files. Furthermore, an intruder with compromised credentials can potentially use that generative AI search tool for **reconnaissance** to more quickly locate valuable data to exfiltrate.

The era of generative AI has also caused changes in terms of **software development workflows**. When an organization seeks to build an AI model today, their starting point is often to adapt a relevant model rather than build their own from scratch. There are a number of public platforms for sharing models and datasets, and while this provides a valuable opportunity to collaborate, it also introduces a new **supply chain risk**. Organizations need to ensure that they are scrutinizing the providence of their code – and appropriately managing the attack surface associated with using **third-party models**. Current DevSecOps methodologies need to be adapted and expanded to encompass the additional attack surface associated with generative AI.

Your Teams are Probably Already Using Generative AI

Public facing, free tools like ChatGPT or coding assistants are very tempting for users. Even if your organization has given guidance or taken measures to block these tools, there is a good chance that some employees are using them on their work or personal devices. This can put **proprietary company data** at risk or introduce a new, unmonitored attack surface. As far as possible, organizations should seek to bring generative AI out of the shadows by substituting sanctioned tools for shadow IT.

This evens applies to the SOC – there's a good chance that a fraction of your team is already experimenting with these tools to accelerate their workflows. It's better to have a **sanctioned strategy** to pursue generative AI use cases and put in place appropriate **governance and guardrails**.

2 | What Are the Top Use Cases for Generative AI in Cybersecurity?

An Approach to Finding Low Hanging Fruit

Generative AI has a wide variety of potential use cases in security. These include large language model (LLM) focused use cases, using conversational assistants, semantic search capabilities, and code generation, as well as multi-modal use cases such as evaluating the risk of an unknown binary. One framework for identifying promising use cases is to find the overlap between the most pressing challenges faced by security professionals and the strengths of generative AI. The highest returns can be found by pursuing use cases that live in this intersection. This has led us to focus on SOC operations – where we see a number of time consuming, manual tasks which are well suited to automation using generative AI capabilities.

Examining the Day-to-day Life of the Analyst

The cybersecurity skilled worker shortage has been a challenge for years, and survey data indicates that the problem is getting worse rather than better. For analysts, this often means that they face pressure in their roles, with many feeling overwhelmed or burned out. Most organizations receive more alerts than they can address, and it is up to analysts to triage and prioritize them, often relying on heuristics to compensate for incomplete context. Most analysts also face the problem of tool overload where they need to constantly jump between tabs to find the information they need.

Similarly, most teams have strategic initiatives that they know should be prioritized, (e.g. ensuring external attack surfaces are hardened) but that advance more slowly than optimal because the team is constantly being torn away to fight fires. When it comes to incident response workflows, most teams have a huge appetite for anything that can help save their analysts time, be that less time processing data, less time clicking between interfaces, or less time writing queries.

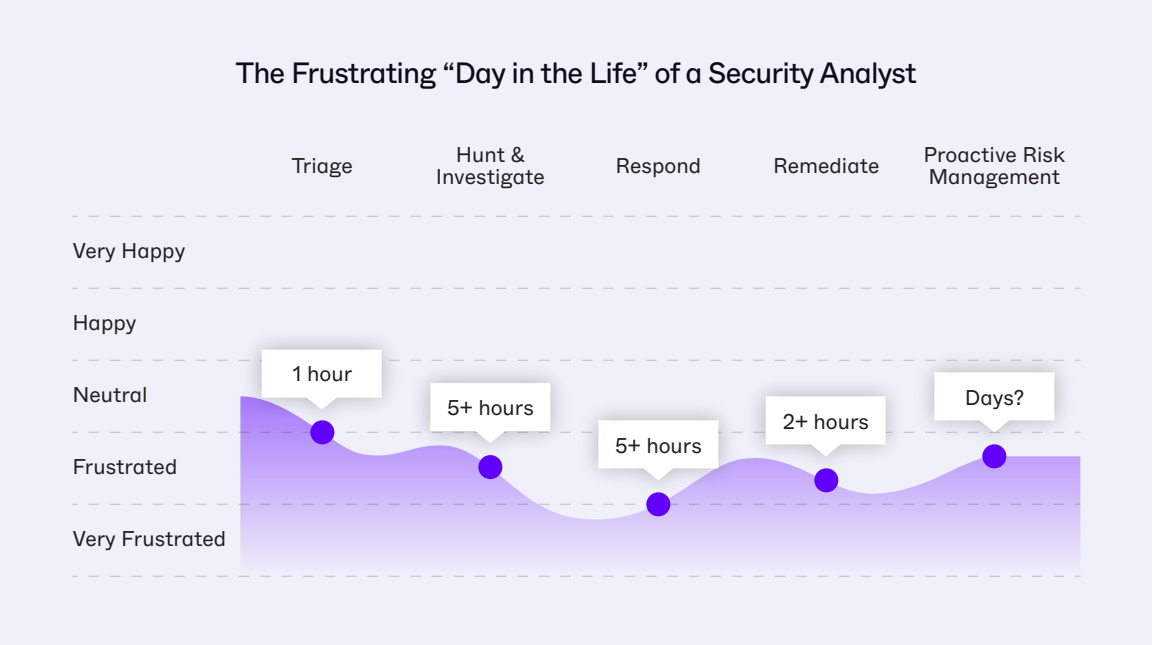

The Strengths of Generative AI

One area where generative AI is very strong is perception; looking across vast quantities of data, correlating related items, and presenting that information in a coherent form for a human to absorb. One significant impact of generative AI on human-machine interaction is that it can enable natural language exchanges, rather than obliging the human to use a coding language or an interface. This works in both directions: not only can generative AI perceive and summarize for a human, it can also understand a natural language prompt to help that person translate intention into action.

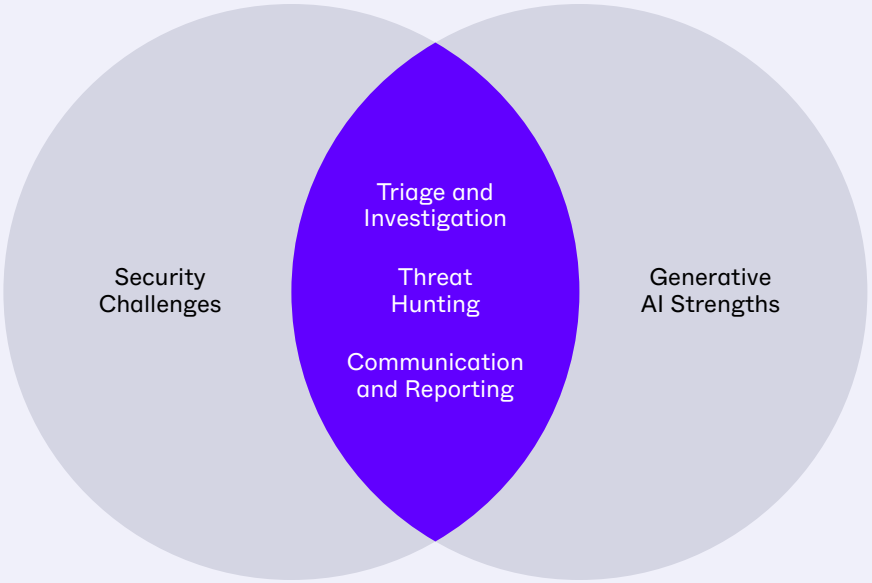

The use cases with the highest returns are those that address the most pressing security challenges and take advantage of generative AI's intrinsic strengths.

Based on this line of analysis, we recommend CISOs consider the following three high-potential use cases:

- **Alert triage and investigation assistance** is often cited as one of the heaviest time burdens, with considerable manual effort to find and analyze relevant logs. Generative AI's search and summarization capabilities present a promising opportunity here.
- **Threat hunting** requires an assessment of complex and dynamic threat intelligence as well as specialized knowledge of coding languages and data schema. Generative AI's ability to translate natural language prompts into code has huge potential to lower the barrier to entry and save time.
- **Coordination and reporting** adds friction when responding to complex incidents, and under routine conditions, takes security professionals away from their core responsibilities. Generative AI can help by automatically recording and summarizing investigation actions, and providing a single view to support response efforts.

Use Case 1: Alert Triage and Investigation Assistance

This use case focuses on the interpretation and prioritization of alerts, as well as the subsequent investigation steps analysts may take if the alert is deemed significant.

Challenges

Today's SecOps team face the challenge of managing a **huge volume of alerts** with scarce personnel. A typical analyst could be reviewing dozens of alerts every hour. Currently, this process begins with triage where a more junior analyst will review a long list of alerts. They may have some degree of prioritization built into their tools or it may be entirely manual. The analyst will use **heuristics** to determine which alerts to examine more closely. The outcome of this step is to reach a verdict on which alerts are highest priority and therefore merit further investigation. The challenges here are both sheer volume, and the fact that the analyst is **working in isolation** reviewing their own alerts, and may lack context collected by other tools or handled by other teams. The result is that triage costs the team valuable time and some risks may be overlooked.

93%

of companies say they cannot address all the alerts they receive on the same day[1]

56%

of large companies handle 1,000+ security alerts each day[1]

The next phase is **investigation**. Here, efforts are complicated as relevant context is spread across different tools. Alerts may involve multiple systems and dependencies, making it necessary to spend valuable time tracing the alert across various components to fully understand its impact. **Integrating and correlating** data from multiple security tools and platforms can be challenging, especially when dealing with proprietary formats. Overall, the result is that a lot of time is spent manually reviewing logs.

Example: An alert is received that many users were deleted at once. The analyst reviews logs to find out which user executed the command. Then, the analyst cross references with network logs to see if data was transferred.

Opportunity

Generative AI has the potential to automate the initial triage phase – so that analysts begin their task with a verdict, rather than working to reach one. It also accelerates investigation by bringing together relevant context that may be spread across different sources, automatically triggering contextual enrichment steps in response to an alert.

[1] 56% of Large Companies Handle 1,000+ Security Alerts Each Day, Dark Reading, 2020

Solution

Generative AI solutions are deployed across thousands of customer environments, and use learnings from this global context to compare a given alert with similar alerts to reach a verdict on whether it is a true positive. Alerts which do not resemble anything the tool is familiar with can be escalated to a human. Depending on the outcome of the initial triage, additional automated tasks are triggered in a chained manner. These tasks enrich telemetry to provide teams with greater context. These steps could include.

- **Contextual Enrichment**
  Gathers additional data and context around the event, such as user accounts, device information, or network traffic patterns.
- **Correlation**
  Determines if the event is part of a broader attack or a standalone incident by correlating it with other events that may have triggered and are stored in a centralized logging backend.
- **Threat Intelligence Lookup**
  Queries threat intelligence feeds to check if the indicators associated with the event are known threats.

This process of chained detections learns from the steps taken by human analysts in similar contexts, and incorporates a feedback loop for continuous improvement.

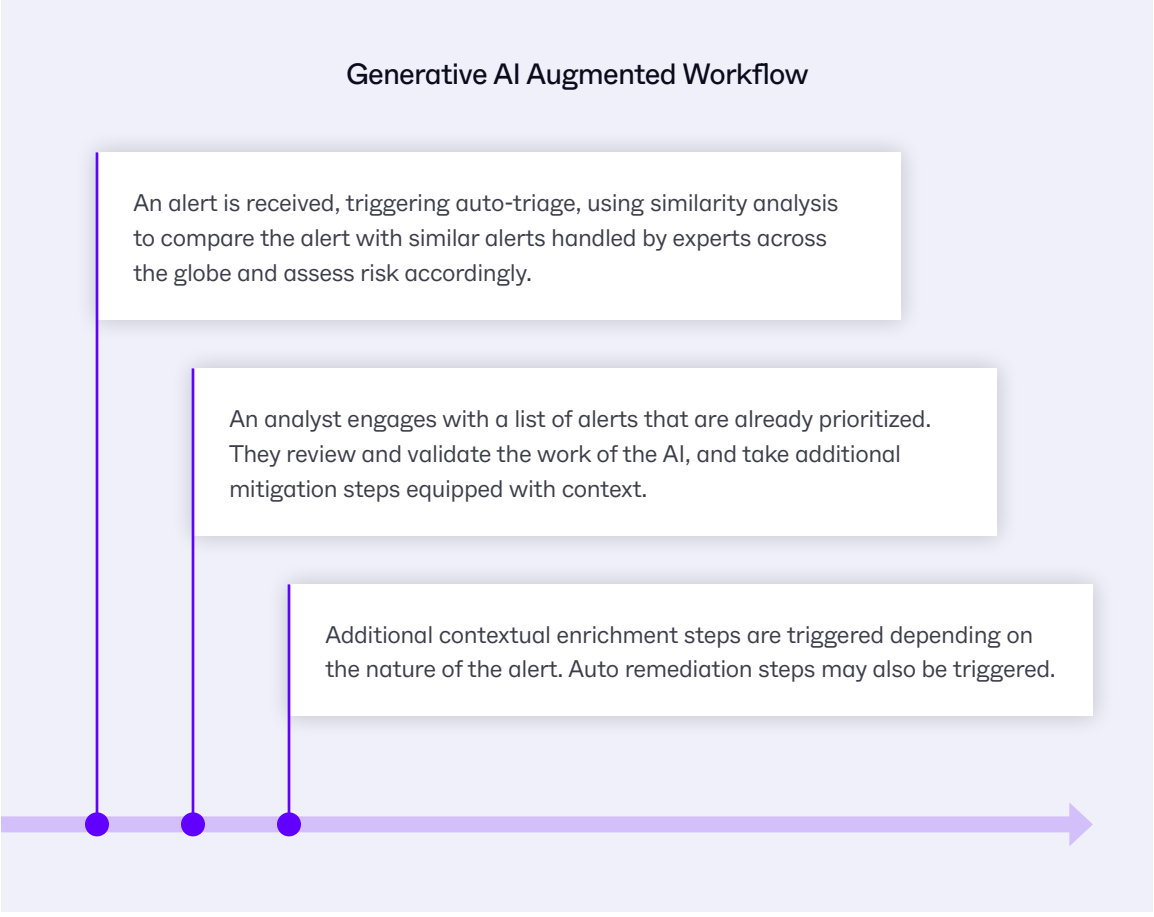

Use Case 2: Threat Hunting

Threat hunting is a security practice aimed at iteratively searching for indicators of compromise. Hunters seek to proactively discover and address threats – as compared with reactively responding to an alert. For many organizations, the time burden and high barrier to entry make it difficult to establish a disciplined formal threat hunting practice.

Challenges

Threat hunting and investigation is the top challenge we hear about when we talk to analysts. It's a long, complex workflow with many highly manual tasks along the way.

The first step for a SecOps team is making sense of threat intelligence and developing a hypothesis about what the most relevant threat might be. As previously mentioned, this has become increasingly challenging, as the **splintering of a dynamic threat landscape** has made it more difficult to focus on key actors when new groups and affiliates surface all the time. For example, a gang may be disrupted by law enforcement one week and reemerge weeks later with new servers and new tactics.

The next challenge is formulating a query. Today, this is a specialized skill that requires an experienced analyst as there are more than twenty threat hunting query languages and none of them are intuitive. Furthermore, formulating an effective query requires an understanding of the **data schema**, so that the right fields and logs are found to answer the question asked. This is a barrier to entry for newcomers to an organization with less familiarity of the data architecture. Even for an expert it is simply not possible to memorize all the facets of the data schema, so looking up fields will cost them valuable time, even if the data is normalized.

What does it take to answer the question "Am I being targeted by a particular threat?"

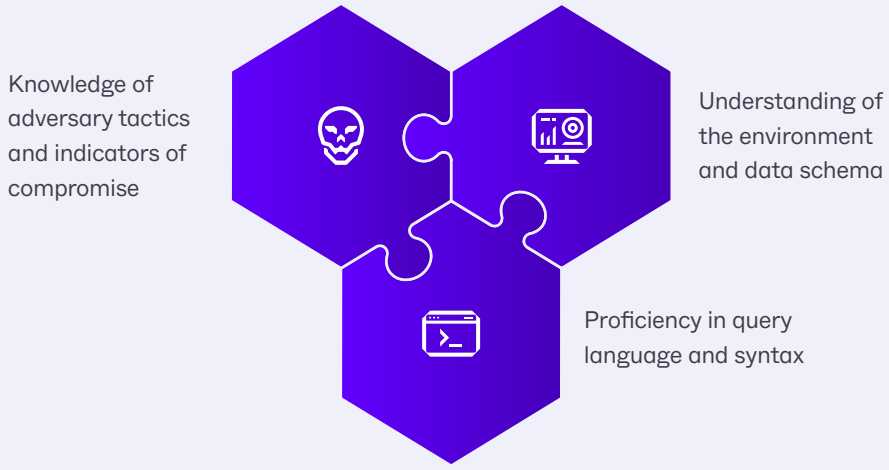

If any element is absent, then threat hunting will not be effective.

Opportunity

Generative AI's capacity to **translate from natural language to code and vice versa** enables analysts to easily and rapidly create queries from a natural language prompt and summarize complex results or threat intelligence. This **lowers the barrier to entry** for threat hunting, empowering junior analysts and helping more experienced threat hunters save time. For more experienced analysts, the biggest time saving is that they **no longer need to manually lookup elements of the data schema**.

Solution

Generative AI can summarize dense threat hunting feeds and answer natural language questions to help analysts focus on relevant areas. Furthermore, generative AI tools can **create pre-prepared queries** based on threat intelligence trends, enabling analysts to start hunting with a **single click**. These can be focused on topics ranging from a particular threat actor, to assets, to tactics, to anomalies, and more.

Rather than write a complex query, analysts can give a natural language prompt. To illustrate the power of this capability, let's explore some example queries:

"SentinelLabs published a report detailing techniques for Lockbit 3.0. Conduct a hunt around the "Event Viewer Tampering" TTP."

```
| filter( event.type == "Behavioral Indicators" AND (indicator.name ==
"EventViewerTampering" OR indicator.name == "EventTampering") ) | group
EventCount = count() by src.process.user | sort -EventCount | limit 1000
```

Generative AI can also then provide a natural language summary of these results, which can be a real time saver for more complex queries:

```
The event viewer tampering has been identified and grouped by username.
The user "jdoe" has the highest event count, followed by "ahughes."
```

From here, generative AI can further augment and accelerate the investigation by suggesting contextual follow-up questions. For example:

- "Which users have accessed the event viewer without proper authorization?"
- "Can you provide a breakdown of event viewer tampering by even type (e.g. process creation, file modification)?"
- "Show me the event viewer tampering activities associated with the user 'jdoe'"

Use Case 3: Coordination and Reporting

Cybersecurity teams are not islands – they are part of a wider organization, and they regularly engage and collaborate with a variety of stakeholders. A response effort to a major incident requires many security professionals – possibly spanning multiple teams – to collaborate closely and to communicate with stakeholders across IT and the broader business.

Challenges

Coordination and communication can be a source of friction and can take precious time teams valuable time even under ordinary conditions. Most organizations want to help their security professionals focus on security rather than writing emails and reduce the burden of **time-consuming, manual reporting tasks**.

In the context of a major incident, these sources of friction can become more acute, as every minute is critical. Key challenges include:

- **Information siloes** with separate tools and teams for different areas of security. For example, some organizations may have separate teams for endpoint security, network security, and identity and access management. It is incumbent on these teams to manually summarize their findings and share with their counterparts.
- **Duplication of effort as team** members lack mutual visibility into one another's efforts.
- **Decision-making and escalations** as complex incidents can cause information overload, forcing team members to spend time ensuring the right people get the right information. Furthermore, in an ambiguous, evolving situation, it can be difficult to determine if circumstances meet escalation criteria, or the appropriate escalation path.

Opportunity

Generative AI can remove much of the manual effort of coordination by automatically assembling relevant information in a single location for all stakeholders to review. It can also leverage this information to save time on common communication tasks by **generating professional drafts of emails and reports**.

When surveyed on the tasks they enjoy the least, security professionals named communicating (email, Slack etc.) as their **number one** most disliked task.[1]

[1] Voice of the SOC, Tines, 2023

Solution

Generative AI brings together a considerable amount of context from different sources as part of alert triage. Throughout the investigation, generative AI records each step taken in chronological order and its results. This information can be made available through **auditable and shareable notebooks**, providing the team with a unified of the investigation. This reduces the burden of manually summarizing and sharing investigation findings and reduces the problem of **duplicated effort**. After the fact, this notebook is also valuable for reviewing the incident and meeting any reporting requirements. Generative AI can draw from this notebook to draft an email or report.

For example: An alert is registered for unusual file access patterns and a spike in CPU usage across several servers. An analyst decides to escalate the alert and quickly generates a summary email. As other team members join the response effort, they read the email and may click into a shared timeline of all the investigation and response steps that have been taken so far.

Looking to the future, we see opportunities for generative AI to support coordination even further by assessing situations and **assigning tasks** to relevant roles accordingly. This could be based on pre-defined rules and playbooks, or it could be something the AI generates dynamically according to context.

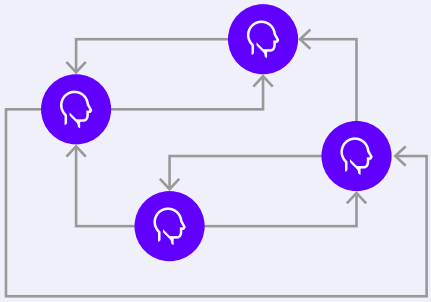

Today, communication and coordination add complexity to a major incident response.

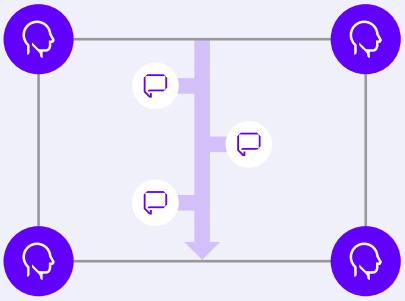

Generative AI brings together context from different sources into a shared view to streamline collaboration.

3 | How Can Organizations Evaluate Generative AI Offerings and Vendors?

Bolted-on Versus Built-in

Given the level of customer interest, most cybersecurity vendors today will claim some generative AI capabilities. **So how do you distinguish marketing hype from the real thing?** The first point to look out for is whether the product is truly integrated with the broader platform or simply a "bolt-on." In the latter case, the vendor is essentially putting a thin layer on top of an LLM without guiding the output or adding additional value to serve security use cases.

Ensuring Accuracy and Preventing Hallucinations

LLMs generate text based on patterns learned from training data. They do not truly understand the content but predict the next word based on statistical relationships. When given vague or contradictory prompts, models might generate creative but inaccurate responses ("hallucinations"). This is why it is important to build an additional layer on top of a basic LLM when tailoring generative AI to a cybersecurity context. Look for the following in a vendor's architecture:

- **Retrieval Augmented Generation (RAG)**
  An approach that combines the strengths of retrieval-based methods and generative models by integrating external knowledge sources.
- **Curated Knowledge Base**
  The vendor's model should draw from a reservoir of cybersecurity context maintained by the solution vendor. When responding to a prompt, the model first retrieves relevant information from this hub, which is then fed into a generative model along with the original query to produce a more accurate and contextually relevant response.
- **Cybersecurity-Specific IP**
  A promising vendor will have a strong track record with traditional AI in a security context, with proprietary insights into threat detection and broad malware datasets to draw from.
- **Meta-Conditions**
  Any generative AI tool oriented to cybersecurity should also be programmed with a set of parameters and constraints that the model must meet when generating a response. These conditions not only prevent "creative" responses, they also help protect the tool from jailbreak or misuse.

The Importance of a Common Data Model

For generative AI to effectively search and summarize across vast datasets, a unified data foundation needs to be in place. A common data model, built around an open standard like the **Open Cybersecurity Schema Framework (OCSF)**, helps realize this full potential by allowing the model to process information from different sources. Furthermore, an open model allows the tool to process data from a variety of vendors, not just the tool's creator.

Protecting Data

Exposing your security logs and telemetry to a generative AI tool carries risk – especially if the data will be exposed to a third-party LLM (such as GPT-4). It is therefore essential to understand how a tool uses your data and ensure that appropriate controls and governance are in place. Factors to consider:

- The security measures in place to protect your data should include encryption, access controls, and compliance with regulations. Verify the vendor's **data handling practices**, including data storage, processing, and retention policies.
- Free generative tools often use data submitted by users to improve the product. This is less likely to be the case for paid enterprise tools but should still be confirmed to ensure **proprietary information** is protected.
- Organizations should understand how **sub-processors** handle their data and confirm that the data is not available to the sub-processor, is not used to improve their product, and is not permanently retained.

Top Questions to Ask Vendors

- Can we see a live interactive demo and try some queries?
- What sub-processor agreements do you have in place?
- What type of regional isolation is used?
- What data protection measures are in place?
- What steps are taken to ensure the privacy of any data we expose to the model?
- How is your product protected against jailbreak or prompt injection?
- How does your architecture reduce the risk of hallucinations and encourage accuracy? Can you share any metrics to quantify accuracy?
- How performant is the model?
- What types of data can the tool answer questions about?
- Is your pricing model transparent and predictable?

4 | What Are the Implications of Generative AI for the Future Role of the SOC??

Today's Limitations Will Not be Tomorrow's Limitations

The opportunity for generative AI is very real – with mature use cases that organizations can adopt today. That said, we haven't yet hit a stable equilibrium and the pace of change is rapid. Our advice to organizations is to think beyond individual use cases and develop a **strategic, forward-thinking approach**. Organizations should seek to adopt a formal process for analyzing workflows, identifying promising use cases, testing new solutions, and ultimately rolling out new capabilities.

We take the bold view that all information security workflows will eventually incorporate generative AI. This technology is more than a new widget; it is a new form of human-machine interaction leading to the emergence of an **autonomous security operations** model. In broad terms, this means automating laborious busywork to enable humans to focus on higher-impact tasks. **AI will automate repetitive, manual tasks, and help humans make sense of complexity, while humans will spend more time validating, deciding, and tackling unusual cases.**

Generative AI Adoption Maturity Model for Cybersecurity

Stage 1: Nascent

- Senior analysts are experimenting with generative AI capabilities.
- Organization is evaluating use cases and solutions.

Stage 2: Mature

- Organization has an enterprise solution integrated into their workflow and trained on their data.
- Organization has formal governance and guardrails around generative AI.
- Organization is actively evaluating and testing additional use cases.

Stage 3: Pervasive

- All infosec workflows are augmented by generative AI.
- Organization has a repeatable process for integrating and extending generative AI capabilities.
- There is a full feedback loop to improve models, encompassing data from the environment and global threat intelligence.

Defining the Appropriate Level of Autonomy

Our perspective is that the next step for the SOC is to move towards "conditional autonomy." This means the system is able to autonomously carry out tasks like threat hunting, vulnerability management, and remediation actions. However, if the system detects anomalies outside its confidence level, it alerts a human analyst to take over. The system has full autonomy under certain conditions, with escalation mechanisms and human readiness to intervene.

An Answer to Cybersecurity's Oldest Problem?

Generative AI has been met with both excitement and hesitancy. There are risks to address, from hallucinations to the impact these tools will have in the hands of threat actors. That said, our perspective is one of cautious optimism. One of the oldest and most pressing challenges for every security team in the world is the volume of tasks and the complexity faced by scarce professionals. **We believe that generative AI will help analysts confront the vast scale and complexity of threats – tipping the balance in favor of defenders.** Crucially, we expect AI automation to deliver a "time dividend" which will enable teams to pivot from fighting fires to spending more time on proactive and strategic activities. For example:

- **Security Campaigns**
  Targeted effort to tackle a security concern, communicating the objectives company-wide and measuring performance. Today most companies will attempt this but these efforts always compete with other priorities. Reducing the burden of incident response frees up time to focus on these strategic campaigns.

- **Exposure and Vulnerability Management**
  Today the sheer volume of vulnerabilities can be overwhelming. Organizations must be strategic about which to prioritize, and many risks will be tolerated for a period of time. AI driven automation, will help free up bandwidth to tackle more of these preventative activities.

- **Proactive Threat Hunting**
  Overall, we see AI shifting the role of analysts form reactive fire fighting to proactive hunting. AI will simultaneously reduce the time burden of incident response while also lowering the barrier to entry for threat hunting.

Alongside benefits like saving time, improving detections, and accelerating investigation and response, it appears that generative AI also has an unexpected benefit. Analysts using these tools seem to enjoy their jobs more – because it allows them to spend less time on mundane, repetitive tasks, and more time thinking critically and problem-solving.

36%

of cybersecurity professionals say that their team spends most of its time addressing high priority or emergency issues and not enough time on strategy or process improvement.[1]

93%

of cybersecurity professionals agree that greater automation would improve their work life balance.[2]

[1] Examining and Addressing Threat Detection and Response Challenges, ESG, 2019

[2] Voice of the SOC, Tines, 2023

Explore How Generative AI Could Unify, Accelerate, and Simplify Your SecOps Workflows Today

Generative AI is already in the hands of attackers. It'salos in use by most defenders, whether via sanctioned tools or through shadow IT. While the risks here are considerable, we are cautiously optimistic. We see this technology transforming cybersecurity by empowering analysts and unlocking time for proactive defense. In the longer term, we see generative AI driving a paradigm shift for SecOp workflows as organizations move towards an autonomous security operations model.

To navigate the risks and begin realizing value from security use cases, organizations need a capable partner. SentinelOne has been a trailblazer in generative AI innovation, both AI-powered security in the Singularity Platform and now generative AI security with Purple AI - the industry's most advanced AI security analyst.

Purple AI's capabilities include:

- Translate natural language into complex threat-hunting queries for search and triage
- Get full visibility with the only generative AI security analyst that supports the Open Cybersecurity Schema Framework (OCSF)
- Summarize event logs and indicators in natural language
- Proactively detect risk with our threat hunting quick starts library and guide analysts with recommended next questions
- Scale collaboration with shared investigation notebooks and email generation

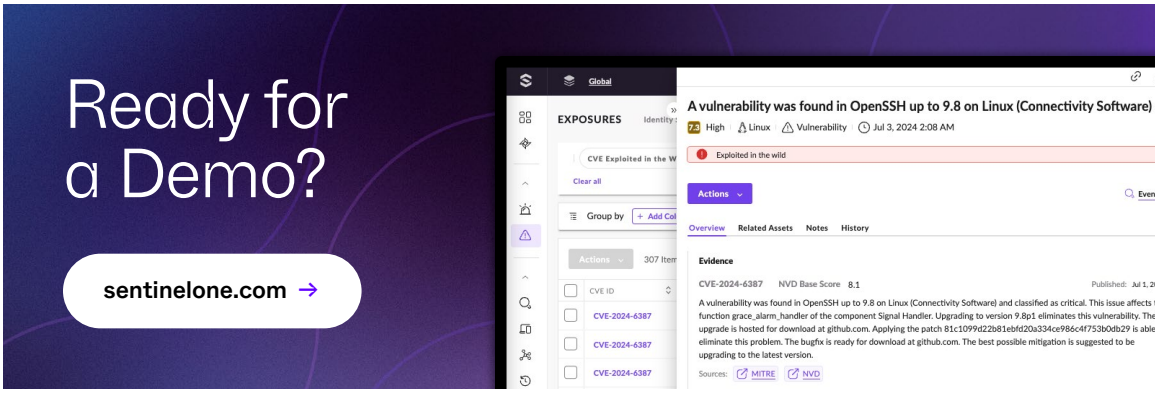

Innovative. Trusted. Recognized.

Gartner

A Leader in the 2024 Magic Quadrant for Endpoint Protection Platforms

MITRE ENGENUITY

Industry-leading ATT&CK Evaluation

+ 100% Detections. 88% Less Noise
+ 100% Real-time with Zero Delays
+ Outstanding Analytic Coverage, 5 Years in a Row

Gartner Peer Insights

96% of Gartner Peer Insights™

EDR Reviewers Recommend SentinelOne Singularity

FedRAMP · TEVORA PCI DSS Attestation HIPAA Attestation · AICPA SOC · STAR LEVEL ONE · 100 VIRUS · SE Labs BEST Innovator WINNER 2021 · AAA · Trusted Cloud Provider CSA

Contact Us

techpartners@sentinelone.com
+1-855-868-3733

sentinelone.com

About SentinelOne

SentinelOne (NYSE:S) is pioneering autonomous cybersecurity to prevent, detect, and respond to cyber attacks faster and with higher accuracy than ever before. Our Singularity XDR platform protects and empowers leading global enterprises with real-time visibility into attack surfaces, cross-platform correlation, and AI-powered response. Achieve more capability with less complexity.

24_MKTG_Product_WhitePaper_013_Four_Questions_CISOs_are_Asking_r3_01072025