Entering the Era of Generative AI-Enabled Security

Introduction

The tidal wave of cybersecurity threats to organizations keeps growing taller. More than 90% of respondents to Deloitte's Future of Cyber 2023 survey reported at least one cyber compromise. Social exploits targeting workers and their networks come alongside attacks on the vast connected infrastructure powering the modern web, enterprise systems, and edge devices.

In a data-driven era, globally connected infrastructures and applications inject cyber risk at each level of an organization's digital activity, and as the power of artificial intelligence (AI) grows, criminals are deriving new ways to automate their tactics. The reality is that the scale of cyber events eclipses the traditional security operations center (SOC) model, where a human analyst reviews and assesses each event. The threat events are simply too numerous and the workforce too taxed to review (much less adjudicate) all of them. As a result, one challenge for today's first line analysis of security events is to single out attacks or compromises from the mass of innocuous data flowing into the SOC from throughout the enterprise's tech stack and network.

By way of analogy, security analysts are attempting to find needles in haystacks when they do not know if a needle exists or what it might look like. What's needed is a way to automate not just the process for finding curious and suspicious activity in the network but also to automate the decision to escalate that finding for human review.

The long-awaited expectation that AI can be used as a force multiplier for cybersecurity is coming to fruition. Deep learning models designed for pattern recognition and predictive analytics are suited to the task of ferreting out cyber threats. Yet, there is an awareness today that off-the-shelf cybersecurity products, often narrowly designed to address a specific use case within the technology stack, are insufficient to meet and keep up with the cyber threat in all its forms.

Recently, a new type of AI has reached the marketplace, and its capabilities are opening the door to new approaches for detecting threats and hardening an organization's cyber posture. The acclaimed generative AI and large language models (LLMs) that have captured public attention for their ability to mimic human speech and comprehension can be used in the cybersecurity environment to overcome challenges in threat monitoring, architecting systems and tools, and talent shortages.

Gen AI capabilities and value for security

The public has become broadly aware of gen AI by way of LLMs that were, in the earliest use cases, sophisticated chatbots. The capacity of an LLM to intake user prompts and output coherent and (usually) accurate replies startled many, and today, there is a growing marketplace of LLMs available for similar applications and public use.

What is perhaps less commonly known is how gen AI is being integrated with tech stacks, both upstream and down, to enhance the function of other systems and the capacity in the workflows. A core gen AI capability is that it enables an intuitive interface with data in real-time. In much the same way as a public user may prompt an LLM to summarize news items, an enterprise can query its own data for insights into processes, product iteration, customer engagement, and indeed, cyber threats. Whereas AI-enabled products for specific use cases require analysts to understand and manage a multitude of tools across an ever-growing attack surface, gen AI solutions can serve as an interpretive bridge across datasets, giving analysts a more natural method for identifying threats and managing enterprise security.

To understand how gen AI can help analysts meet today's cyber threat landscape, consider the areas where a security LLM can enhance capacity while improving threat detection and remediation.

Threat monitoring at scale

Cyber threats are overwhelming in volume and novelty. As one attack avenue is closed, another is cracked open. One challenge is that cyber criminals may covertly penetrate a system and then lurk there, for months, in advance of a future data exfiltration or network attack. In some cases, the compromise may not even be noticed until the attackers are on their way out with troves of sensitive information.

What's needed is a way to rapidly identify potential threats, analyze them, and present the insights in a consumable way allowing analysts to make decisions and take action. The vast majority of data exchanges and network connections are legitimate and unthreatening, making it challenging to sift through the noise and find the vulnerabilities and suspicious data in a meaningful and impactful timeframe. Rules-based approaches will only find instances of known attacks; novel attacks may go unseen. Machine learning is suited to finding patterns or anomalies in data, allowing security analysts to leverage AI to identify the hard-to-find threats. And yet, this requires a high degree of technical proficiency and investment, from training and testing the model to managing its function over time.

Enter: Gen AI.

An LLM trained and fine-tuned for security can supercharge an organization's existing cyber platforms and software, helping analysts sift through information and identify potential threats that deviate from expected, legitimate activity. With prompts, the LLM can classify, synthesize, and summarize these insights in an intuitive, digestible way and in a variety of formats (e.g., text, data visualizations).

A true next-generation security LLM doesn't stop there. It can work across silos of data, querying enterprise systems and insights from third-party providers to reveal threats that might otherwise go unseen. Threat intelligence combined with point-in-time incident analysis empowers security analysts to identify and contain threats before they spread.

The toil of architecting systems and tools

One of the challenges for cybersecurity professionals is managing the multitude of tools available for specialized use cases. The cybersecurity market has responded to a proliferation of threats with a multitude of products, but the problem is that more security tools provide more security alerts, which requires more security analysts with specialized skills to respond. The toil involved in the process of reviewing so much specific, highly technical data strains capacity and can complicate the process of threat mitigation.

It's said if everything is a priority, nothing is a priority. Gen AI can be leveraged to break through this logjam. The adoption and integration of a security-specific LLM can bolster an organization's cyber capacity and elevate it to a new level by alleviating some of the process and toil associated with summarizing and prioritizing response to threats.

With gen AI and a cloud service provider (CSP), organizations can reach across datasets and connect information from enterprise assets, the CSP's threat information, public reports, and expert security research and analysis. When reports and structured data sources can be summarized with a prompt and Tier 1 security alerts can be triaged with automation, security analysts are empowered to focus their knowledge and effort on higher priority threats. The result is a reduction in needless toil and a corresponding increase in targeted cybersecurity capabilities.

Overcoming the talent gap

Given the vast threat landscape and the digitization and data modernization taking place across every industry, it is no surprise that cybersecurity talent remains in high demand and short supply. The talent shortfall means not only that organizations are challenged to do more with less, but it also means that many people who are responsible for cybersecurity operations did not train as security specialists. Developers, administrators, engineers, and junior analysts are called to take up the cybersecurity mantle, even as the complexity and sophistication of cyber threats and tools only compounds.

According to a Deloitte survey, even high-cyber maturity organizations cite a lack of skilled cyber professionals as a key challenge in managing cyber risk. Industry trends do not suggest the talent gap will be closed any time soon, and as a result, organizations require creative and workable approaches to maximize cybersecurity capacity in the absence of sufficient human talent. One viable approach could be empowering more people from across the enterprise to participate in threat identification and security operations. Gen AI and its ability to consult and present complex data in conversational, consumable language is ideally suited to helping non-security professionals contribute to threat monitoring and reduce toil for security professionals.

A security-specific LLM can help make security operations accessible to the entire organization. By pairing the workforce with the intuitive capabilities of an LLM that can access a variety of datasets and threat intelligence, the enterprise can begin to address the Tier 1 challenge of sorting through the hay to ferret out the needles. The result is that in-demand security professionals can operate "at the top of their license," focusing on the most urgent or priority threats that pose the greatest risks to the organization.

Longer term, the use of LLMs may fundamentally change how cybersecurity is accomplished. Assistive AI enabled by gen AI both upstream and downstream creates a force multiplier effect, in which more people can take part in cyber operations with greater access to information and analysis, yielding a lower net burden on security professionals.

Gen AI for cybersecurity with Deloitte and Google Cloud

Using LLMs to address core challenges in cybersecurity is a brave new frontier, which Google is tackling today. The Google Cloud Security AI Workbench, powered by Google's security-specific model, is a platform for adding gen AI functionality to security products and is based on years of foundational AI research by Google. It is designed to help address the core challenges limiting cybersecurity operations today: the scope and scale of the threats, the toil of architecting security tools, and the stubborn talent gap.

Preparing the enterprise to use Security AI Workbench presents organizations with a great opportunity to address some key security challenges. With its Google Cloud Security Specialization, Deloitte is recognized as having a strong proficiency and experience with securing Google Cloud environments. Our services around cloud transformation, AI governance, and cyber and strategic risk services allow us to deliver end-to-end cloud security solutions and transformational cloud-native SOC programs. We are named as a leader in The Forrester Wave™: Cybersecurity Incident Response Services, Q1 2022; managed cloud security services (by IDC); and strategic risk management consulting (by ALM).

Just as important, we bring a leading approach to AI risk management and governance. Deloitte's Trustworthy AI Framework™ helps organizations develop effective AI governance and promote compliance by assessing models and use cases across dimensions of trust and ethics. Leveraging the framework to enable governance, including for gen AI, Deloitte helps clients across the professional services of strategy, advisory, and audit and assurance, bringing to bear our breadth of human capital and depth of market-leading domain knowledge to help companies meet today's cybersecurity challenges. For organizations thinking through concerns specific to AI systems, such as poisoning of the training data, injecting malicious inputs through prompts, and extracting confidential information in the training data, Google designed the Secure AI Framework, inspired by industry best practices, to help assess and mitigate risks.

A new and growing landscape of threats demands a new approach to managing cybersecurity. With the power of a gen AI-enabled platform, enterprises have a pathway to boost their workforce, prepare for the threats that will emerge in the future, and better secure their organizations.

Contact us

Mike Morris
Managing Director
Deloitte & Touche LLP
micmorris@deloitte.com

Arun Perinkolam
Principal
Deloitte & Touche LLP
aperinkolam@deloitte.com

Jacob Crisp
Global Head of Strategic Response
Google Cloud
Learn more

This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited ("DTTL"), its global network of member firms or their related entities (collectively, the "Deloitte organization") is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication. DTTL and each of its member firms, and their related entities, are legally separate and independent entities. Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited ("DTTL"), its global network of member firms, and their related entities(collectively, the "Deloitte organization"). DTTL (also referred to as "Deloitte Global") and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about learn more.