aws marketplace

GENERATIVE AI SECURITY EBOOK

The Generative AI Revolution and What It Means for Cybersecurity

By SANS Digital Forensics and Incident Response Instructor
Matt Bromiley

In collaboration with:

SANS

Sponsored by AWS

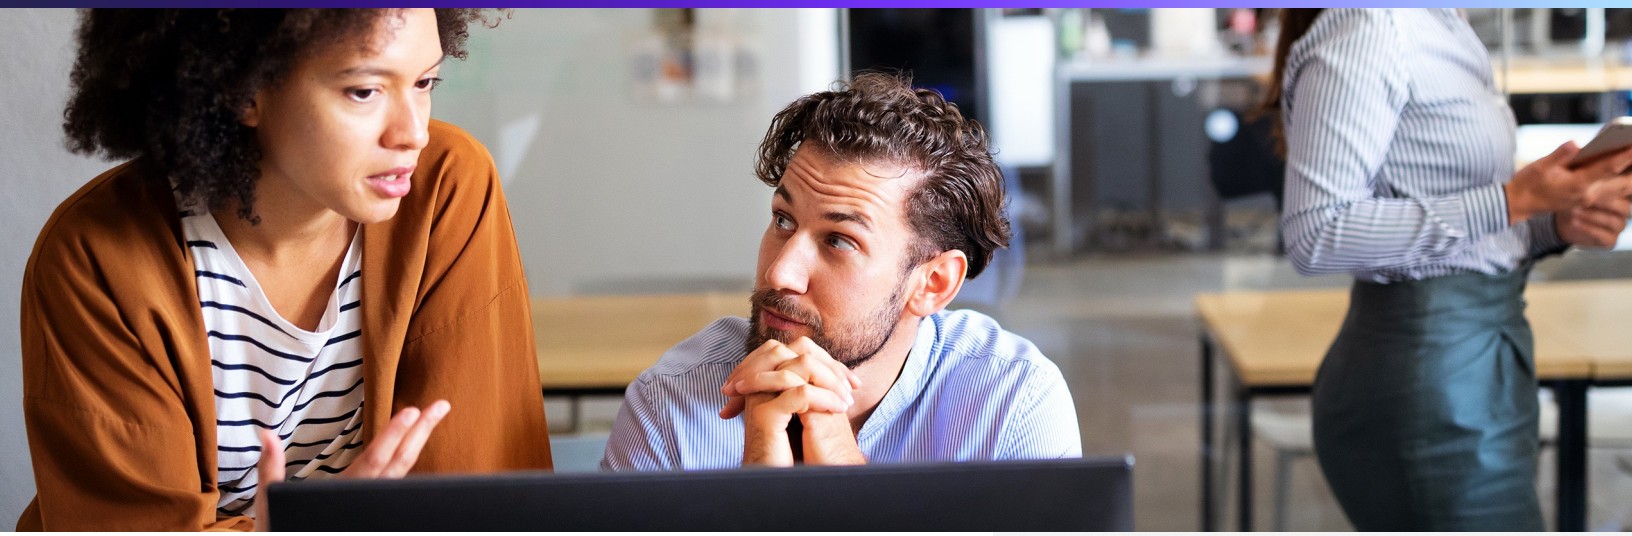

Introduction

In the *Generative AI for Security: Harnessing AI for Customer Impact* webinar, AJ Yawn, SANS Institute associate instructor and partner in charge, product and innovation at Armanino LLP, and Mark Weiss, strategic initiatives leader, Amazon Web Services (AWS), provide a helpful roadmap for securely harnessing generative AI for development and data governance with Amazon Bedrock. You can watch it here on demand.

In addition to the webinar, you will want to read this ebook, in which SANS digital forensics and incident response instructor Matt Bromiley examines key ethical considerations and technical challenges around generative AI.

The author will explain how to account for skill gaps, how to determine organizational readiness, and ways that generative AI can benefit your business.

The solutions for this use case can be found in AWS Marketplace:

Trellix

elastic

LaunchDarkly

Browse AWS Marketplace to discover these and other products that enhance your overall cloud security posture.

Learn more by visiting AWS Marketplace →

SANS | Research Program

eBook

The Generative AI Revolution and What It Means for Cybersecurity

Written by **Matt Bromiley**

August 2024

aws marketplace

Generative AI Is Everywhere ...

... and security is struggling to keep up! The explosive growth of generative AI (GenAI)—and the breakneck adoption across a wide array of business use cases—is presenting a challenge to enterprise security teams. However, GenAI poses significant new challenges while simultaneously creating enormous opportunities. This double-edged sword has the potential for an increase in security events due to AI-generated vulnerabilities, while simultaneously being used to detect advanced threats.

Unfortunately, GenAI remains poorly understood in many areas and the implications for enterprise security teams are no exception. Many security organizations do not yet grasp just how GenAI implementations will impact them or the enterprises, users, and data they protect. This knowledge gap unfortunately results in lack of preparedness for new technologies, process changes, and skill sets required to manage GenAI effectively.

To help security teams combat this, Amazon Web Services (AWS), in collaboration with the SANS Institute, has developed a duo of e-books focused on GenAI in security. This first offers strategic, high-level insights into the most critical challenges and benefits this new technological paradigm brings. The second, takes a more tactical approach, presenting detailed guidance for enhancing GenAI security, real-world use cases, and showcasing GenAI products and services available in AWS Marketplace.

Explore: Understanding the GenAI Landscape

We begin our journey by exploring the **rapid expansion of GenAI**. These technologies and capabilities have exploded onto the scene, mostly driven by industry forces and advances in machine learning (ML) and natural language processing (NLP). The ability to generate human-like text, images, code, videos, and much more has opened a vast array of business applications. For many organizations, these have arrived in the form of chatbots or "co-pilots," but also may include things like automated content creation, sales processes, and even production-close code authoring.

Like any new technology or capability, a range of opportunities and challenges arise. We'll examine both below.

Opportunities

The GenAI landscape provides a wealth of opportunities for organizations, security teams, customer management and automation, and many more. A few notable opportunities include:

- **Enhanced threat detection—**GenAI can assist with analyzing vast datasets to identify potential security threats quicker and more accurately than traditional methods.
- **Automation of repetitive tasks—**GenAI can automate routine security tasks, such as log analysis, ticket management, evidence generation, and components of incident response. This frees up security personnel to focus on more strategic activities.

- **Improved decision making with GenAI—**With deeper insights and predictive analytics, GenAI can enhance and add more confidence to security operations decision making.
- **Broader coverage and increased scalability—**GenAI applications can deliver real-time threat detection capabilities and 24/7 coverage, reducing the need for human hours—especially on lower-level tasks.

Challenges

Although the opportunities are vast, GenAI does not come without its own set of challenges. These can include:

- **Data security and privacy—**The use of large datasets in training GenAI models can expose sensitive information if not managed properly.
- **Regulatory compliance—**It's easy to simply "turn GenAI on," especially if offered as a feature or part of an existing toolkit. However, ensuring that GenAI applications comply with data protection requirements (such as GDPR or PCI-DSS) can be complex. Organizations must implement robust data governance frameworks, which may add more strain to already-taxed security teams.
- **Model transparency and "explainability"—**GenAI models can be complex and opaque, which makes it difficult to understand their decision-making process. Model transparency is essential to building trust and understanding the "process" through which decisions are made.

GenAI Misuse

Despite its short history, the potential risks of careless or unmanaged GenAI use have already surfaced. In February 2024, an open-access scientific journal published an article with "gibberish descriptions and diagrams of anatomically incorrect [mammals]."[1]

The Knowledge Gap in Enterprise Security

It's becoming easier and easier for security teams to "enable" GenAI within their tech stacks—particularly by vendor implementations. However, many security teams still grapple with understanding how GenAI will impact their operations. This knowledge gap can leave them vulnerable to new types of threats and unprepared for rapid technological changes. Addressing this gap will require:

- **Education and training—**Security teams need ongoing training to stay current with the latest GenAI technologies *within their environment* and the associated security implications. This includes understanding new attack vectors and learning how to conversely leverage GenAI for defensive purposes.
- **Innovation and adaptability—**Organizations must foster a culture that encourages innovation with new technologies while maintaining a strong focus on security. GenAI use cases are no exception. Although a development team may find enormous value and speed by utilizing GenAI-generated code, it is not guaranteed to follow regulatory or business-specific requirements.

[1] "Science journal retracts peer-reviewed article containing AI generated 'nonsensical' images," February 16, 2024, https://venturebeat.com/ai/science-journal-retracts-peer-reviewed-article-containing-ai-generated-nonsensical-images

Develop: Building a Secure GenAI Framework

Building a secure GenAI framework should be a critical first step, not an afterthought. However, despite the ease of "enabling" GenAI capabilities, there are important considerations when developing a model that is intended to replicate the actions of humans or confidently inform human decision making. Some careful considerations include:

Bias and Fairness

Mitigating bias in AI systems is critical. Bias can stem from multiple areas, including (but not limited to):

- **Data bias—**This type of bias occurs when training data is not representative of the larger population (sampling bias) or there are inaccuracies in the data-collection process (measurement bias). Bias can inaccurately sway decisions for GenAI-related detection and response.
- **Algorithmic bias—**Bias can be reinforced through feedback loops, where biased decisions lead to biased data, and so on.
- **Cultural or societal bias—**This type of bias can reflect cultural assumptions or historical bias, depending on the data and models used to train. For example, detection rules written to only observe one language or dialect can incorrectly categorize other language content.
- **Anchoring bias—**Bias that stems from disproportionate influences in the training data might overly rely on a pattern even if new data suggests otherwise.

Data Privacy

Protecting sensitive data used in GenAI decision making and algorithm training is paramount. This can include things like business intellectual property (IP), personally identifiable information (PII), and personal health information (PHI). GenAI models should be designed to handle data securely and to comply with privacy regulations. Critical features can include:

- Encryption and access controls can protect data from unauthorized access. Data anonymization techniques can be used to enhance privacy and ensure data is not leaked via a GenAI implementation.
- Data governance policies ensure data is handled responsibly through its life cycle. This includes the full path of data flows, from collection to storage, processing, and sharing.

Skills Gaps and Organizational Readiness

One of the biggest concerns for any organization is going to be: "Once we have GenAI implemented, are our teams ready to work with it?" Enterprises need to identify the skill sets required to manage GenAI effectively. This may involve hiring new personnel with specific expertise in AI and ML, as well as offering training to upskill current staff.

GenAI implementations also can require changes in organizational processes and workflows. Security needs *must* be "at the table" to discuss and enhance the broader enterprise strategy so that GenAI implementations are developed and deployed securely. Ensuring **cross-functional collaboration** between security teams and other departments can ensure that security considerations are front of mind.

Ensuring AI System Integrity

When exploring the implementation of a GenAI application or incorporation via a current tech stack, it is imperative to protect these systems from manipulation, tampering, or data poisoning. We discussed biases earlier. In this situation, we are discussing intentional poisoning or tampering with the goal of achieving a bias or influencing decisions or results. Basic implementations can include data encryption and *strict* access controls.

GenAI systems, especially custom-designed systems, are no different from any other application. They require a strong focus on security. Although the threats against AI may still be in their infancy, adversaries are developing their tradecraft. AI-focused **threat modeling** can help identify potential vulnerabilities in GenAI systems. This takes a proactive approach, monitoring the threat landscape for potential GenAI weaknesses, adversary exploits, and/or tactics, techniques, and procedures (TTPs). Appropriate security measures should be implemented.

It goes without saying that **continuous monitoring** around AI systems can help organizations detect and respond to suspicious activities in real time. We won't go into the full depths of creating a security monitoring approach in this e-book, however, it's valuable to understand that GenAI systems need protection just like other systems.

Deploy: Integrating GenAI Securely

Deploying GenAI applications requires blueprinting and architecture diagrams to illustrate how GenAI services and tools integrate into a secure deployment. AWS offers a suite of tools that can work together to help protect against some of the data privacy concerns we discussed above. These include:

- **Amazon Bedrock—**Amazon Bedrock provides a secure and scalable platform for developing and deploying GenAI applications. Its features include data encryption, access controls, and integration with other AWS security services.
- **AWS PrivateLink—**AWS PrivateLink enables secure and private communication between GenAI applications and other AWS services. This reduces the risk of data exposure and ensures compliance with privacy regulations.
- **Amazon S3—**Amazon S3 offers secure storage for GenAI training data and models. Its features include encryption, versioning, and access controls to protect sensitive information.
- **AWS Key Management Service (KMS)—**AWS KMS provides encryption key management for GenAI applications. This ensures that data is encrypted at rest and in transit, enhancing security.

Use Cases

Implementing GenAI applications or taking advantage of available add-on implementations has already proven successful for many organizations. Using GenAI applications can provide a huge boost to productivity and/or customer satisfaction, especially for lower-level or easily automated tasks.

Use Case: Healthcare

A healthcare organization used GenAI to analyze patient data and predict disease outbreaks. Obviously, this involves extremely sensitive and potentially identifiable data, such as PII and PHI. Along with tools mentioned above, Amazon Macie is a security service that can assist in discovering and protecting sensitive data types. This can be critical in ensuring that only the *right* data is consumed and utilized by a GenAI application.

The healthcare industry has fallen victim to many types of cyberattacks over the past few years. With the GenAI and security capabilities of AWS, the organization was able to securely deploy applications that analyzed patient data with minimal security risks. Furthermore, the use of capabilities, such as Amazon GuardDuty, offers continuous monitoring and threat detection capabilities.

Use Case: Financial Services Organizations

Financial services organizations have two burdens GenAI can assist with. First, financial institutions are frequent sources of advanced cyberattacks. GenAI can be an invaluable tool to detect suspicious or intrusive activity from both internal and external parties. With the right data and trained models, GenAI can help detect attacks and mitigate or automate remediation activities.

Second, financial organizations also mine their own data looking for fraudulent transactions and account activity. This is another excellent use case for GenAI, as applications can be trained to recognize risky patterns in account or transactional data and facilitate automated actions. With AWS security tools, designing and implementing both detections and fraudulent data mining can be done in a secure manner.

Continuous Monitoring and Feedback

As mentioned above, continuous monitoring of GenAI systems is essential to ensure data and model integrity and application effectiveness, in addition to the needs of security teams. Key monitoring capabilities include:

- **Establishing and monitoring key performance metrics for GenAI systems—**This can help track effectiveness and accuracy.
- **Monitoring for anomalous behavior—**Security and development teams should monitor for anomalous behavior, of both the model and the system(s) around the model. Cloud-hosted resources, even if service-specific, can still fall suspect to attack.

Finally, establish feedback loops for ongoing improvement of GenAI systems. This includes gathering input and data from security teams, end users, developers, and other stakeholders to hone GenAI applications and/or use cases over time. Feedback can include **user feedback** to provide insight into systems' performance and to make iterative improvements. In addition, implementing a culture of **continuous improvement** ensures that GenAI uses grow and adapt to match the needs of the organization *while* meeting the requirements of the security team.

Closing Thoughts

The transformative power of GenAI offers significant opportunities for enhancing enterprise security. However, it also presents new challenges that require careful consideration, proactive management, and a consistent security-first culture. By understanding the GenAI landscape and developing a secure framework for adoption and integration, enterprises and security teams can harness the benefits of GenAI technology while mitigating risks.

The next e-book will provide step-by-step guidance, real-world case studies, and a look at the broad range of GenAI security solutions offered by AWS Marketplace. With these insights and recommendations, security teams can effectively navigate the complexities of GenAI and strengthen their security posture without weakening their attack surface.

Sponsor

SANS would like to thank this paper's sponsor:

aws marketplace

aws marketplace

Why use AWS Marketplace?

AWS Marketplace is a digital software catalog that makes it easy to find, try, buy, deploy, and manage software that runs on AWS. AWS Marketplace has a broad and deep selection of security solutions offered by hundreds of independent software vendors, spanning infrastructure security, logging and monitoring, identity and access control, data protection, and more.

Customers can launch pre-configured solutions in just a few clicks in both Amazon Machine Image (AMI) formats and SaaS subscriptions, with entitlement options such as hourly, monthly, annual, and multi-year contracts.

AWS Marketplace is supported by a global team of solutions architects, product specialists, and other experts to help IT teams connect with the tools and resources needed to streamline migration journeys to AWS.

How to get started with generative AI solutions in AWS Marketplace

Watch the webinar:

Generative AI for Security: Harnessing AI for Customer Impact

Watch on demand →

Discover solutions:

Start exploring and experimenting with solutions such as Amazon Bedrock.

Visit AWS Marketplace →

Talk to an expert:

Speak to a solution architect who can help solve your business challenges.

Get connected →

In collaboration with

SANS