mission cloud

EBOOK

EBOOK

Safeguarding the Future with GenAI Security

In an era dominated by data, the rise of generative artificial intelligence (genAI) is paving the way for unprecedented business opportunities and substantial challenges.

GenAI, fueled by sophisticated language models, is a powerful catalyst for innovation and optimization. It promises to bring a revolution in operational efficiency and new product development. However, this revolutionary technology has risks, like any other innovation.

Business leaders must navigate the complex landscape presented by genAI — ensuring their organizations can harness its potential while simultaneously implementing robust measures to mitigate associated security risks.

This guide is your strategic blueprint for understanding and reinforcing genAI security. It will equip your business with the knowledge and tools required to secure and prosper in this transformative period, allowing your organization to flourish and adapt in a world where technology and innovation are paramount.

The GenAI Landscape

Companies are already using genAI to generate compelling text, images, and videos. The technology, while still evolving, is receiving immense interest and substantial investments. It's pivotal for companies to meticulously navigate the prevailing hype cycle, discerning genAI's current standing to manage expectations and foster progress.

"When we think about where we are with genAI, we're somewhere at the beginning of this structure," said Jarret Raim, Mission Cloud's chief information security officer.

He notes that a recent PwC survey found that nearly all business leaders say their company is prioritizing at least one initiative related to AI systems in the near term, but only 35% of executives say their company will do the following:

- Improve AI systems and process governance
- Ensure AI-driven decisions are interpretable and easily explainable
- Monitor and report on AI model performance
- Adequately protect AI systems from cyberthreats and manipulations.

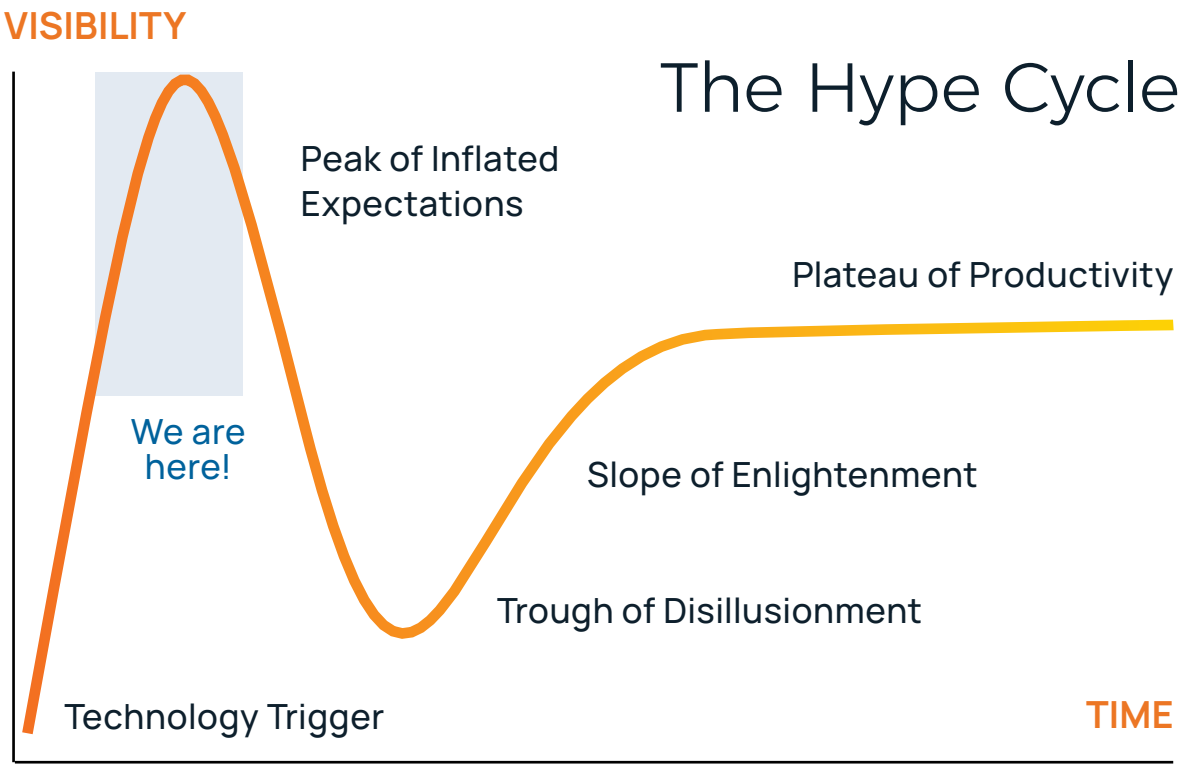

That said, genAI's unparalleled creative capabilities have attracted considerable attention, including investments from private equity and venture capital. For the first half of 2023, investment in generative AI startups received record equity funding — topping $14.1 billion across 86 deals, according to CB Insights.

Navigating genAI's evolution requires a comprehensive understanding of its position on the hype cycle. This journey from heightened anticipation to profound understanding is crucial for effectively managing expectations and directing progress. As companies delve deeper into the capabilities of genAI, they must balance enthusiasm and realism — aligning ambitions with the practicalities and potentials of the technology.

"The jury's out here. We don't know what it's going to be useful yet," Raim said. "So we want to enable ourselves to go and try those things. And how do we do that while also maintaining the commitments that we've made to ourselves and to our customers?"

Amid the fervor surrounding genAI, companies must gain clarity about this technology's realistic applications and limitations. GenAI providers, meanwhile, are catching up to the needs of the business community.

"This market segment is very nascent," Raim said. "A lot of these providers don't actually even offer the ability for us to say, 'Hey, we want to be able to upload data to your model and have that model use that data. But we don't want that data available to anyone else. We just want it to be available solely to us.' That's coming. It's going to happen."

GenAI Security Fundamentals

GenAI could add value to the plethora of data that remains uncollected. As companies increase their use of genAI, safeguarding even less common data becomes critical, as every piece of information has the potential to be transformative.

4 Underlying GenAI Risks

GenAI models are essentially stochastic engines. That means they generate output based on probabilities, and often 'hallucinate' if not prompted well. Business leaders must understand the four key risks of genAI:

DATA RISKS

Error propagation, data leakage, lack of legal approvals to use the data for AI models, and misleading and harmful content caused by low-quality data used in the models create many data risks.

MODEL AND BIAS RISKS

Model hallucination producing information that is factually incorrect, AI models containing biases that lead to discriminatory or unfair outputs can produce legal and operational risks for businesses.

PROMPT OR INPUT RISKS

Users who input misleading, inaccurate or harmful information into the AI models can be a source of risk that companies must address.

USER RISKS

Relying on genAI in circumstances where users lack the expertise to check its accuracy can produce misinformation and harm.

"These models are really good at providing you information that looks very compelling, but many times, it's incorrect," Raim said. He noted that, in his experiments with genAI, the technology creates code that looks like it could work but fails to function.

The pillars for adopting genAI are fundamentally rooted in data privacy and trust. The integrity and confidentiality of data aren't just operational necessities; they're competitive advantages. Establishing trust-centric practices around data is essential in fostering an environment where genAI can flourish responsibly and ethically.

GenAI Security in Practice

Integrating genAI within business operations requires a comprehensive plan to manage risk.
Raim outlined a 10-step path forward he uses at Mission Cloud.

01 **ESTABLISH** a governance structure and enterprise-wide generative AI risk management framework.

02 **ENGAGE** with leading AI technology providers.

03 **PERFORM** legal diligence on your contracts, model performance, and intellectual property.

04 **ENGAGE** your employees in identifying use cases in their work and for the company's customers.

05 **EVALUATE** and prioritize use cases based on risk and reward. Look for common patterns that apply to the majority of use cases, are reusable and applicable to future use cases. Start with those use cases.

06 **BUILD** your generative AI factory, with tooling and enhancements and with the appropriate security and controls.

07 **DO** trial runs in sprints.

08 **ROLL OUT** genAI for broader use under a dedicated enterprise program office.

09 **MONITOR** your foundation models and applications for compliance and drift periodically. Use a governance tool to do this.

10 **ADOPT** robust generative AI systems and model metrics and monitor for concept drift, toxicity, hallucination and bias.

Taking the first step may be daunting given the shifting nature of genAI. However, Raim shared how Mission Cloud's genAI policy can be simplified into three main points:

- → Only share data with companies where the company has a contractual relationship.
- → All generated content is reviewed, approved and owned by a qualified human.
- → All automatic usage is logged and audited for reproducibility.

One principle that should guide any AI policy: "Don't take humans out of the loop," Raim said. "We won't use AI to replace humans, we'll use it as an augmented tool to make us better at our jobs."

Creating a GenAI Security Policy

In the dynamic landscape of artificial intelligence (AI), one issue towers above the rest — data privacy. It's the bedrock of security in the AI era and demands immediate attention. Here are examples of activities that companies should consider approving or prohibiting when crafting their genAI security policies:

APPROVED

- Asking an AI model about public tools and APIs
- Refining public text, such as marketing descriptions of products or templates that aren't specific to customers
- Running local AI models on your laptop or private server
- Using genAI tools, such as Amazon Bedrock or CodeWhisper, in company accounts to build and design new models.

PROHIBITED

- Uploading source code in development to an AI tool that the company doesn't have a legal agreement with
- Providing a spreadsheet of customer monitoring data and asking the AI model to summarize or process that data
- Asking the model to review or improve a communication being sent to a specific customer
- Generating images using AI platforms for use in public or customer-facing use cases because of the lack of clear copyright ownership

GenAI continues to evolve, and enterprises must change with it. But the use of AI models in security is still new, and it's unclear how their role will evolve. What's crucial right now is keeping even the smallest amounts of proprietary data out of AI models when the company doesn't have a contractual relationship with the tool's provider.

"Once information enters the model, the model knows it," said Jonathan LaCour, Mission Cloud's chief technology officer. "It will probably not spit that out if you ask a generic question, but a bad actor could ask a specific question. And if that specific information is available to the model, it will do its job and repeat it back to you."

The era of genAI brings forth a paradoxical challenge — while AI models bolster our capabilities, they also provide malicious actors with new tools.

Staying Ahead of GenAI Threats

Companies need to confront the ever-changing concerns and challenges surrounding genAI security. That is why it's vital to work with an experienced partner, especially one experienced with Amazon Web Services (AWS).

Mission Cloud can help you explore genAI's potential while fortifying against future uncertainties. We helped real estate platform provider myTheo use machine learning and genAI to increase revenue by offering greater value to the company's customers without requiring additional locations or facilities. With the help of MAP funding through AWS, Mission Cloud developed a proof of concept for three new revenue-generating solutions.

Reach out to schedule a complimentary session with a solutions architect to explore how Mission Cloud can augment your company's genAI security, shielding against the unforeseen and fostering a secure environment for innovation to thrive.

mission cloud

Mission Cloud is a leading AWS Premier Tier Services Partner and Cloud Managed Services Provider. Through its dedicated team of expert cloud operations professionals, cloud analysts, and solutions architects, Mission Cloud delivers a comprehensive and differentiated suite of agile cloud services designed to help businesses migrate, manage, modernize and optimize their AWS cloud environments.

1 (855) 647-7466 • www.missioncloud.com • sales@missioncloud.com

aws PARTNER Premier Tier Services

- Amazon Redshift Delivery
- Amazon QuickSight Delivery
- Data & Analytics Services Competency