Deloitte.

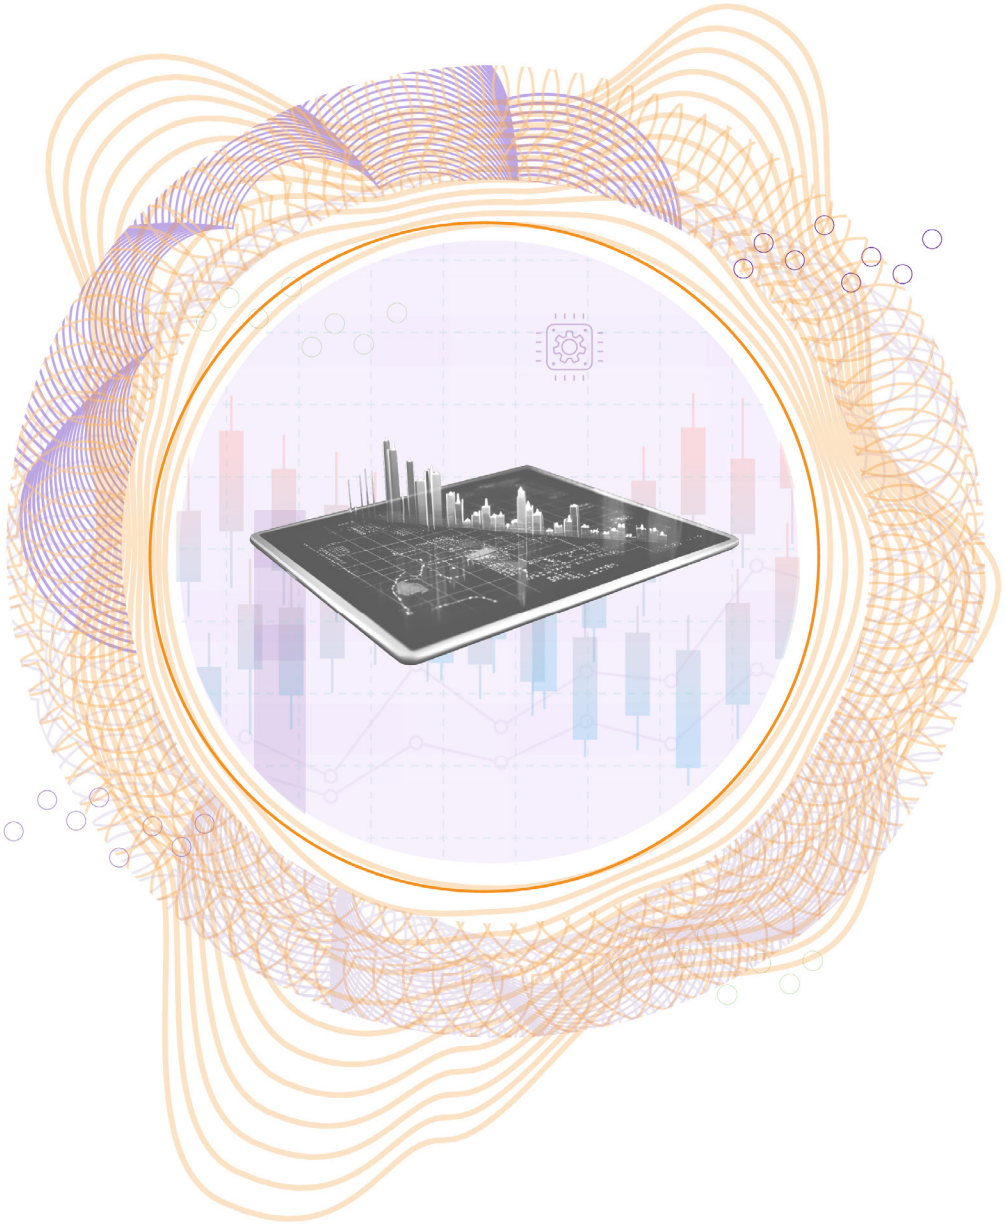

Navigating the impact of generative AI on security

A CISO's guide

Evolving the role of CISO with the advent of Gen AI

Cyber has been ranked as the most important risk globally[1] for the second year in succession. Reports have shown that average costs from such incidents reached an all-time high in 2022 and will continue to increase at a multi-fold pace in the coming years. The role of the chief information security officer (CISO) will likely assume an even greater strategic significance within the organization's cybersecurity program. The insurance industry, in particular, is being targeted by a myriad of cyberattacks as it possesses a great deal of personally identifiable information (PII) and protected health information. Research has found that customer and employee PII is the costliest to have compromised at $183 per record.[2] As insurance companies migrate toward digital channels to create tighter customer relationships and offer new products, a new wave of investment is directed toward advanced analytics and generative artificial intelligence (Gen AI). Although these investments provide new strategic capabilities, they also introduce new cyber risks and attack vectors to organizations. The challenges are likely to become more complex as insurers prepare to leverage large language models (LLMs) for Gen AI, which will require not only collecting and handling large amounts of sensitive data, but also safely exposing it across multiple applications, interfaces, and cloud platforms.

As business and product teams find new ways to use AI, organizations must also ensure its safe, secure, and ethical use. With growing digital disruption to conduct business, and diminishing trust, the need for a skilled CISO has grown manyfold. The CISO is responsible for managing security across a distributed network to ensure that the data remains secure; maintaining compliance with regulatory requirements; and educating employees and informing executives about cybersecurity risks (82% of the largest insurance carriers have been the focus of ransomware attacks from cybercriminals[3]). They not only have to contend with AI's immediate impact, but also prepare for how it will shape their responsibilities in the future. Highlighted below are some of the ways Gen AI could add to the responsibilities of a CISO:

- **Data security and privacy:** Assess how models handle sensitive data and ensure they comply with data protection laws and regulations.
- **Access control:** Implement robust access controls to ensure that only authorized individuals have access to systems.
- **Model integrity and security:** Protect AI models from tampering and reverse engineering. This includes ensuring that the models themselves are securely stored.
- **Logging and monitoring:** Establish logging and monitoring systems to detect and respond to security incidents.
- **Training and awareness:** Provide training and raise awareness among employees and stakeholders.

Staying informed is the first step. A CISO should make time to be curious and continuously learn about new developments and how they can affect insurers' security posture. The CISO role will likely evolve from being the "de facto" accountable person for treating cyber risks to being responsible for ensuring business leaders have the capabilities and knowledge required to make informed, high-quality risk decisions.

Potential benefits of Gen AI

Gen AI has the potential to add contextual awareness and decision-making to enterprise workflows and can radically change how we do business. The far-reaching impacts and potential value when deploying Gen AI are accelerating experimental, consumer, and (soon) enterprise use cases. According to Gartner, 68% of executives believe that the benefits of Gen AI outweigh the costs, compared with just 5% who feel the risks outweigh the benefits.[4] The two big advantages of Gen AI today are its capacity to process huge amounts of data at high speed and its ability to communicate clearly. Still, questions remain about how insurance carriers can use Gen AI to bring effectiveness, efficiency, and understanding, while managing the associated cybersecurity risks.

Gen AI security use cases

- **CISO executive reporting:** Gen AI can help streamline and automate the process of drafting reports, including requirements for incident response, threat intelligence, risk assessments, audits, and regulatory compliance. It can provide real-time insights into an organization's risk profile, including its threat landscape, risk levels against critical vulnerabilities, current cybersecurity posture, compliance requirements, and cybersecurity performance metrics, which can all be of aid to insurance CISOs.[5]
- **Act as security assistant:** A Gen AI security assistant can assist security analysts in sifting through piles of log entries to evaluate possible security threats by providing an assessment in seconds. With a single prompt, Gen AI can scour logs and other data and report back on what may be an immediate threat and what isn't. It can also explain and add valuable context to the threat identifiers.
- **Act as application development assistant:** Gen AI can also act as a secure application development assistant. Many code generation tools are embedding security features, and application security tools are already leveraging LLM applications that can help in common security use cases such as vulnerability detection, false-positive reduction, and mitigation suggestions.
- **Mitigate social engineering attacks:** Insurance companies also run the risk of losing money due to whaling attacks (a type of social engineering attack where cybercriminals send executives a spoof email to dupe them into authorizing massive cash transfers). LLMs not only generate text, but they are also helpful in detecting and watermarking AI-generated text, which could become a common function of email protection software. Identifying AI-generated text in social engineering attacks can help to detect phishing emails and polymorphic code.
- **Alleviate security talent and skill shortage:** The sheer amount and complexity of data and threats have become increasingly difficult to tackle. The integration of Gen AI into several security operations tools enables cybersecurity teams to scale while remaining lean and focused. This new interface can reduce the skill requirements for using the tool, shorten the learning curve, and allow more users to benefit.

Apart from the previously mentioned use cases, our first paper in the series delves into further instances across the insurance value chain where Gen AI is utilized, exploring its implications. You can find more details here: Implications of Gen AI for insurance.

Key risks of Gen AI on CISOs' watchlist/threat landscape

Despite numerous benefits, Gen AI was one of the top concerns among security executives over the first few months of 2023.[6] CISOs may feel pressure to allow use of Gen AI broadly, but doing so indiscriminately could create unreasonable risk.

Risks associated with Gen AI for a CISO at the enterprise level generally stem from:

- **Data and privacy confidentiality:** Enterprise use of Gen AI may result in access and processing of sensitive information, intellectual property, source code, trade secrets, and other data, through direct user input or an application programming interface (API). Sending confidential and private data outside of the organization's servers could trigger legal and compliance exposure, as well as risks of information exposure. Such exposure can result from contractual (e.g., with customers) or regulatory obligations (e.g., CCPA, GDPR, HIPAA, CPP Model law) that are in place and relevant to the organization.
  - *Mitigating measure:* Adhere to relevant regulations, such as GDPR or CCPA, to help safeguard sensitive information and maintain customer trust, and use secure Gen AI platforms.
- **Data poisoning/prompt injections:** Corrupt/polluted (poisoned) data leads to malicious or unintended outcomes and can affect the accuracy and reliability of the LLM. By using carefully designed inputs, attackers can manipulate LLMs, compelling them to carry out the attacker's desires. This manipulation can occur by directly altering the system prompt or manipulating external inputs, which may result in serious issues like data exfiltration.
  - *Mitigating measure:* Implement robust access controls to ensure that only authorized individuals can access systems.
- **Enterprise, SaaS, and third-party security:** Due to Gen AI's wide adoption and proliferation of integrations, there are concerns that data would be shared with third parties at a much higher frequency than earlier anticipated, posing a threat to non-public enterprise data and third-party software. For example, third-party applications leveraging a Gen AI API, if compromised, could potentially provide access to email and the web browser, and allow an attacker to take actions on behalf of a user.
  - *Mitigating measure:* Establish comprehensive data governance policies and procedures. This includes defining data ownership, data classification, and data life cycle management. Clearly define who has access to AI-generated data, how it is stored, and for how long. Additionally, organizations must implement data quality controls and establish mechanisms for data lineage and audit trails. By adopting a robust data governance framework, enterprises can help mitigate risks associated with Gen AI.

In addition, a CISO should be consulted and informed about the following risks:

- **Legal and regulatory risk:** Legal and compliance risks arise from the fact that the legal and regulatory landscape surrounding Gen AI is still nascent. Consequently, enterprises may not be aware of all the legal requirements they need to comply with when using this technology. When Gen AI is used as part of a regulated use case in consumer-facing communications, whether for direct consumer interactions or to produce consumer-facing materials (such as consumer information notices), regulatory or private law may include requirements and create liability.
  - *Mitigating measure:* Comply with relevant data protection regulations, and refrain from sharing customers' sensitive information and the organization's own sensitive data. Consider a platform that operates inside the secure network of the organization. Obtain explicit user consent when collecting and using personal data for Gen AI purposes.
- **Bias and discrimination:** Training on biased data may lead to illegal discrimination, potential damage to reputation, and possible legal repercussions for the enterprise as Gen AI may not be aware of potentially defamatory, discriminatory, or illegal content.
  - *Mitigating measure:* Ensure that the training data used for Gen AI models is diverse, representative, and free from biases. Regular monitoring and auditing of the models' outputs can help identify and address any potential biases, promoting fairness and inclusivity in AI-generated content.
- **Copyright and ownership/risk to intellectual property (IP) rights:** Gen AI models are trained on diverse data, which might include copyrighted and proprietary material, raising ownership and licensing concerns between the enterprise and other data sources used for training.
  - *Mitigating measure:* CISOs should seek a tool that operates end to end on their company's network and does not require users to send sensitive data to external servers or third parties. CISOs should also consider collaborating closely with legal teams to establish robust IP protection measures.

Gen AI security: Focus areas for CISOs

A CISO should ask questions and provide guidance to help leaders create an organizational AI strategy. A comprehensive AI strategy provides guidelines for its usage and factors in legal, ethical, and operational considerations. If used responsibly and with proper governance, Gen AI can provide businesses with many benefits across automated processes and optimized solutions. A comprehensive AI strategy can help ensure privacy, security, and compliance. It should consider the following key questions:

- Who is using the technology in the organization, and for what purpose?
- How can I protect enterprise information (data) when employees are interacting with Gen AI? Do we have governance and contingency in place (i.e., usage and controls)?
- How can I manage the security risks of the underlying technology? How do I balance the security trade-offs with the value the technology offers?

Deloitte's approach to responsible AI[7]—Trustworthy AI™—delivers trust by design throughout the AI life cycle. It's relevant to executives at every level:

- The CEO and board set the strategy with special attention to public policy developments and to corporate purpose and values.
- Chief risk and compliance officers oversee control, including governance, compliance, and risk management.
- Chief information and information security officers take the lead on responsible practices, such as cybersecurity, privacy, and performance.
- Data scientists and business domain specialists apply responsible core practices as they develop use cases, formulate problems and prompts, and validate and monitor outputs.

How an insurer intends to use Gen AI and its impacts should be thoroughly assessed across the following five key areas before embarking on a Gen AI journey:

1. **Strategic considerations**
    - Impact of data and AI: Consider the moral implication of uses of data and AI and codify them into your organization's values.
    - External policy and regulation: Understand public policy and regulatory trends to align compliance processes.
2. **Internal controls**
    - AI governance: Enable oversight of systems across the three lines of defense.
    - Internal compliance: Comply with organization policies and industry standards.
3. **External regulations[8]**
    - Cross-industry: Comply with applicable external regulations. For example, recently enacted EU AI Act 2023[9] is a comprehensive guide to AI law with clearly defined transparency requirements and risk levels. The EU AI Act classifies insurance as a high-risk industry, which leads to more stringent regulations and greater transparency industrywide, and firms have to ensure higher system compliance levels to prevent any penalties.
    - Insurance-specific: Focus on insurance regulations being established to safeguard insurance-specific risks. For example, Colorado's draft AI regulation 2023[10] guides life insurers' use of external consumer data and information sources. It outlines requirements that ensure usage of algorithms and predictive models (i.e., AI models) do not result in unfairly discriminatory insurance practices with respect to race.
    - The National Association of Insurance Commissioners (NAIC) has outlined a draft bulletin[11] that provides guidelines for insurers to use while utilizing AI systems (AIS) and ensuring compliance. It emphasizes the importance of AIS programs, AI governance, and documentation.
4. **Risk management**
    - Fair/not biased: Define and measure fairness and test systems against standards.
    - Transparent and explainable: Enable transparent model decision-making.
    - Responsible and accountable: Use policies to clearly establish accountability for AI outputs.
    - Robust and reliable: Enable high-performing and reliable systems.
    - Privacy: Develop systems that preserve data privacy.
    - Safe and secure: Design and test systems to prevent data harms.
    - Role-based access control: Implement robust authentication and authorization mechanisms to restrict access to sensitive Gen AI systems and data.
5. **Leading practices**
    - Use-case identification: Identify the concrete problem you are solving for and whether it needs an AI or machine learning solution.
    - Industry standards: Follow industry standards and best practices.
    - Continuous monitoring: Implement continuous monitoring to identify drift and risks.

Path forward to balance Gen AI's challenges and opportunities

Gen AI offers immense potential for innovation and creativity across the insurance value chain and processes. In this journey, it is critical for the office of the CISO to tackle the unique security challenges and ethical issues and stay on top of the ever-changing regulations in this domain. With business teams eager to leverage Gen AI at scale as early as possible, here are a few change management guardrails to consider in the short and medium term to mitigate risks to insurers' security posture:

- Educate employees on the potential risks of Gen AI usage through in-person training, online courses, and awareness workshops.
- Communicate the importance of transparency and accountability to prevent bias, hallucinations, and other risks.
- Identify and protect sensitive training data, enforce access controls, and implement data loss prevention to prevent leaks.
- Establish clear usage policies, assessment frameworks, and diligence models to evaluate the credibility of third-party AI solutions, with the do's and don'ts of using AI-generated content within the organization.
- Form an approval board comprising stakeholders from different business units to define internal policies based on a risk assessment framework, and oversee adherence to the same while implementing Gen AI use cases.

Contacts

Sandee Suhrada
Principal
Deloitte Consulting LLP
ssuhrada@deloitte.com

Vishvam Raval
Senior consultant
Deloitte Consulting LLP
viraval@deloitte.com

Sunny Aziz
Principal
Deloitte & Touche LLP
saziz@deloitte.com

Sharat Viswanathan
Senior consultant
Deloitte Consulting LLP
sharviswanathan@deloitte.com

Rohan Shinde
Manager
Deloitte Consulting LLP
roshinde@deloitte.com

Meenakshi Rawat
Senior consultant
Deloitte Consulting LLP
merawat@deloitte.com

Endnotes

1. Allianz, *Allianz risk barometer 2024*, January 2024.
2. IBM, *Cost of a data breach report*, 2023.
3. Eliot Partnership, "The evolving role of chief information security officers in the insurance industry," April 17, 2023.
4. Gartner, "Gartner poll finds 45% of executives say ChatGPT has prompted an increase in AI investment," press release, May 3, 2023.
5. Michael Sentonas, "Introducing Charlotte AI, CrowdStrike's generative AI security analyst: Ushering in the future of AI-powered cybersecurity," CrowdStrike, May 30, 2023.
6. Deloitte, Trustworthy AI™, accessed January 23, 2024.
7. Gartner, "Gartner survey shows generative AI has become an emerging risk for enterprises," press release, August 8, 2023.
8. While we acknowledge that regulations pertaining to Gen AI are continuously evolving, we have outlined some of the newly drafted regulations as of October 1, 2023, for reference.
9. European Parliament, "EU AI Act: First regulation on artificial intelligence," updated December 19, 2023.
10. Colorado Department of Regulatory Agencies, Division of Insurance, SB21-169 – Protecting Consumers from Unfair Discrimination in Insurance Practices, accessed January 23, 2024.
11. National Association of Insurance Commissioners (NAIC), "NAIC Model Bulletin: Use of Algorithms, Predictive Models, and Artificial Intelligence Systems by Insurers," exposure draft, July 17, 2023.

Deloitte.

About Deloitte
Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee, and its network of member firms, each of which is a legally separate and independent entity. Please see www.deloitte.com/about for a detailed description of the legal structure of Deloitte Touche Tohmatsu Limited and its member firms. Please see www.deloitte.com/us/about for a detailed description of the legal structure of Deloitte LLP and its subsidiaries. Certain services may not be available to attest clients under the rules and regulations of public accounting.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional adviser. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Copyright © 2024 Deloitte Development LLC. All rights reserved.
8242799