



# Auditing in the Age of Generative AI

April 2024

**CAQ**

# About the Center for Audit Quality

The Center for Audit Quality (CAQ) is a nonpartisan public policy organization serving as the voice of U.S. public company auditors and matters related to the audits of public companies. The CAQ promotes high-quality performance by U.S. public company auditors; convenes capital market stakeholders to advance the discussion of critical issues affecting audit quality, U.S. public company reporting, and investor trust in the capital markets; and using independent research and analyses, champions policies and standards that bolster and support the effectiveness and responsiveness of U.S. public company auditors and audits to dynamic market conditions.

Please note that this publication is intended as general information and should not be relied on as being definitive or all-inclusive. As with all other CAQ resources, this publication is not authoritative, and readers are urged to refer to relevant rules and standards. If legal advice or other expert assistance is required, the services of a competent professional should be sought. The CAQ makes no representations, warranties, or guarantees about, and assumes no responsibility for, the content or application of the material contained herein. The CAQ expressly disclaims all liability for any damages arising out of the use of, reference to, or reliance on this material. This publication does not represent an official position of the CAQ, its board, or its members.

# Contents

4	Introduction
5	Overview of GenAI: What Auditors Need to Know
9	The Regulatory Environment
11	Considerations When Auditing Companies That Are Deploying GenAI
16	Example Use Cases
18	Additional Audit Considerations
18	Conclusion

# Introduction



Collective interest in and awareness of generative AI (genAI) has grown exponentially since the public release of several genAI chatbots powered by large language models beginning in November 2022. While artificial intelligence (AI) and machine learning are not new, the accessibility and ease of use provided by genAI chatbots and similar large language models have led to increased use by individuals and companies. A recent CAQ survey found that one in three audit partners see companies in their primary industry sector deploying or planning to deploy AI in their financial reporting process.<sup>1</sup> This number will likely continue to grow as companies explore the ways in which AI, including genAI, can streamline or enhance accounting and financial reporting operations and processes.

This publication explores some fundamental principles of genAI, new risks arising from its use in processes relevant to financial reporting (financial reporting processes) or internal control over financial reporting (ICFR), and related audit implications. Although some of the considerations discussed may also be applicable for other types of AI, the focus of this publication is specifically on genAI.

One in three audit partners see companies in their primary industry sector deploying or planning to deploy AI in their financial reporting process.

**CAQ's Audit Partner Pulse Survey, Fall 2023**

<sup>1</sup> TheCAQ.org | Audit Partner Pulse Survey | Fall 2023

# Overview of GenAI: What Auditors Need to Know

In order for auditors to identify where and how companies are using genAI in financial reporting processes and ICFR and the risks that could arise from its use that may be relevant to the audit, it will be helpful to have a foundational understanding of some fundamental principles of genAI, including key features of the technology and how it differs from other technologies that companies may be using. As the genAI technology, use cases, and regulatory environment are rapidly changing, it is important for auditors to continue to monitor developments.

## WHERE DOES GENAI FIT WITH OTHER AI TECHNOLOGIES?


AI includes a broad range of technologies, of which genAI is a subset. While there are other types of AI beyond those shown to the right, this graphic depicts where genAI fits with other categories of AI technologies.

**Artificial Intelligence** | AI broadly refers to machines that mimic human-like cognitive abilities. AI includes capabilities such as natural language processing, problem-solving, pattern recognition, anomaly identification, and decision-making. An example of AI is an online language translation service.

**Machine Learning** | Machine learning is a subset of AI that uses algorithms to learn from and make predictions or decisions based on data. Machine learning algorithms are designed to learn and improve from experience. Machine learning is useful for identifying patterns, extracting insights, and making informed predictions. Different methods of machine learning include supervised learning, unsupervised learning, and reinforcement learning.<sup>2</sup> An example of machine learning is a system used by a streaming service that provides recommendations to customers based on their viewing habits.

**Deep Learning** | Deep learning is a subset of machine learning that uses algorithms that roughly approximate the structure and capabilities of the human brain. Deep learning algorithms can simulate an array of neurons in an artificial neural network that learns from vast sources of data enabling the technology to handle complex tasks similar to how humans can. An example of deep learning is driverless cars which can recognize and respond to different situations on the road.

**GenAI** | GenAI refers to a subset of deep learning based on probabilistic technology that can create content, including text, images, audio, or video, when prompted by a user. GenAI creates responses using algorithms that are often trained on open-source information, such as text and images from the internet.<sup>3</sup> Through its ease-of-use, genAI has democratized artificial intelligence making the technology accessible to any user, whereas other types of artificial intelligence have generally only been accessible to data scientists. AI chatbots, like ChatGPT and Copilot, are examples of genAI.



<sup>2</sup> For further discussion of these methods of machine learning, refer to the AICPA and CPA Canada's *A CPA's Introduction to AI: From Algorithms to Deep Learning, What You Need to Know* publication.

<sup>3</sup> Science & Tech Spotlight: Generative AI | U.S. GAO



## HOW DOES GENAI WORK?

### Learning and Generating New Content

GenAI technologies are trained on large datasets where they learn patterns, structures, and representations from the training data. For example, based on the training dataset, genAI learns grammar and syntax and uses its advanced predictive capabilities to mimic knowledge on a wide range of topics. Based on this training data, when prompted by a user, genAI technologies make predictions of the next character, word, phrase, pixel, etc. to formulate a probable response to the user prompt.<sup>4</sup>

GenAI technologies are predictive technologies, and therefore, the outputs are based on what the genAI technology has determined is a probable response. If a user asks the same question multiple times, they might get different answers each time. Different answers may result because genAI technologies are designed to generate varied responses and are trained on diverse datasets, which leads to a wide range of probable responses to a single prompt.<sup>5</sup> Accordingly, genAI technologies are especially helpful for tasks that need creativity or diversity of responses, including generating new content or information, but genAI may not always provide reliable or repeatable information. GenAI technologies do not work like search engines finding facts within their training data but are instead creating new coherent, human-like text.

### Foundation Models and GenAI Technologies Supported by Those Models

When developing and deploying genAI technologies, companies may build and train their own models,<sup>6</sup> or they may begin with a foundation model. Foundation models are large language models that can be adapted to a wide range of downstream tasks, providing the basis for various genAI technologies.<sup>7</sup> There are many foundation models currently available. One example is GPT-4, which is the foundation model used by one version of ChatGPT. This same foundation model can also be the basis for other applications. For example, a company could also use GPT-4 as the basis for its own internal chatbot.

Companies can build their own customizations on top of foundation models. Customizations may include incremental training with the company's own data and fine-tuning the model for specific uses within the company. Using a foundation model can allow companies to develop custom genAI technologies without the significant effort involved in developing their own model. However, companies using foundation models may not have visibility into the data and methods used to train the foundation model.

<sup>4</sup> Prompts are the information (such as a question, command, etc.) entered into a genAI technology to generate a response.

<sup>5</sup> It is possible to configure certain genAI technologies to provide more deterministic responses (i.e., provide consistent and predictable responses). However, the diversity of the datasets that genAI technologies are trained on will still lead to a range (albeit narrower) of probable responses to a prompt.

<sup>6</sup> Although it is possible, it may be rare for companies to build and train their own genAI large language models.

<sup>7</sup> [Explainer: What is a foundation model? | Ada Lovelace Institute](#)

<sup>8</sup> AI hallucination is a phenomenon wherein a large language model (such as a genAI chatbot) perceives patterns or objects that are nonexistent or imperceptible to human observers, creating outputs that are nonsensical or altogether inaccurate. See further discussion at [What are AI hallucinations? | IBM](#).

### CONSIDERATIONS FOR AUDITORS

The probabilistic nature of genAI is a key distinction from other technologies that auditors may have historically encountered in a company's financial reporting processes, which may inform auditors' identification and assessment of risks of material misstatement, including the identification of process level risks or risks arising from IT. Further, when performing audit procedures over information generated by genAI, auditors need to be aware that the information produced by genAI is not necessarily factual and may not be able to be replicated by the same genAI technology, even if the same input is provided again, which may influence how auditors design and execute audit procedures.<sup>8</sup> Auditors' responsibility to obtain sufficient and appropriate audit evidence under applicable auditing standards remains unchanged.

### CONSIDERATIONS FOR AUDITORS

A company may develop its own model, build customizations on top of a foundation model, or use a pre-built solution based on a foundation model (such as a publicly available chatbot) depending on their specific needs. The risks arising from the company's use of genAI will vary depending on the nature of the genAI technology. Auditors may consider the following questions:

- + Did the company develop its own genAI model or is the genAI technology built on a foundation model?
- + If the company is using a foundation model, which foundation model supports the genAI technology?
- + Did the foundation model require incremental training or customization to support the company's use case?

## Explainability and Interpretability of GenAI

There is an increasing desire for genAI users to understand how and why the technology arrives at certain conclusions, which relates to the explainability and interpretability of genAI. Explainability refers to explaining or understanding the underlying mechanisms in the genAI technology's behavior – in other words, *how* the technology made the decision.<sup>9</sup> Interpretability refers to when humans can readily understand the output of the genAI technology through the reasoning behind predictions and decisions made – in other words, *why* the technology made the decision.<sup>10</sup>

A challenge of AI is that it can be a “black box,” meaning that the process to arrive at a specific output is not readily explainable or interpretable, resulting from the inherent complexity of AI algorithms and the nonlinearity of the relationships between the underlying data and the outputs or decisions made. While this challenge exists for all types of AI, including genAI, explainability and interpretability needs will vary depending on a number of factors, including the level of reliance on the technology (i.e., whether the technology is used to augment work performed by an employee or replacing the employee), the nature or type of the output (i.e., whether the output can be independently replicated by a human reviewer), and the level of human in the loop involvement (see further discussion in the *Responding to Identified Risks* section). Additionally, the ability to explain and interpret outputs may be impacted by whether the technology is built on a foundation model or a model developed by the company (i.e., whether the company controls the underlying algorithms). These factors are important for auditors to consider how the use of genAI technologies impacts the company's financial reporting processes or ICFR and the related audit response in tests of controls or substantive procedures.

Explainable AI (XAI) is an emerging area of research focused on techniques to enhance the explainability and interpretability of AI (including genAI). Some of these techniques include embedding features that can provide information regarding the AI technology's confidence in its outputs or decisions or to document the key elements of the input that the AI technology focused on to make its decision. While embedding such features may not be feasible for existing technologies, particularly those genAI technologies built on a foundation model, it may be possible to add certain features on top of genAI foundation models to enhance explainability and interpretability.

## WHY AND HOW ARE COMPANIES DEPLOYING GENAI?

Companies are noting significant opportunities from deploying genAI, particularly from using genAI to enable knowledge workers to perform their jobs more efficiently and effectively. GenAI can help employees streamline certain activities such as those that involve drafting content, summarizing data, and working with unstructured data, among others, which frees them up to focus on more challenging, analytical, or higher-risk tasks. Further, genAI can uncover trends, patterns, and anomalies in large amounts of data that would otherwise be difficult or time-consuming for human employees to uncover manually.

## CONSIDERATIONS FOR AUDITORS

The impact of the black box concept on the audit generally depends on the factors described to the left. Effective human oversight to address explainability and interpretability risks becomes important specifically as companies place heavier reliance on genAI technologies, use cases in financial reporting processes and ICFR become more sophisticated, and outputs from the technology are unable to be independently replicated. The following questions may be helpful for auditors to consider:

- + Is the company placing reliance on genAI technology to generate outputs that are not, or cannot be, verified or replicated by employees?
- + If the company is placing reliance on employees to review the output from genAI technology, how has the company determined that employees have the appropriate knowledge and skills to do so?
- + If the company is placing reliance on genAI technology, how has the company determined that the genAI technology is sufficiently explainable and interpretable for the intended use?

<sup>9</sup> Artificial Intelligence Risk Management Framework (AI RMF 1.0) (nist.gov)

<sup>10</sup> Ibid.

Generally, companies deploying genAI within financial reporting processes will initially use it to augment processes (rather than fully automate them), which enables efficiency but does not eliminate human judgment and decision-making. Particularly in financial reporting processes and ICFR, humans continue to be involved to oversee, understand, and evaluate the relevance and reliability of the outputs from genAI technology. In the future, companies may evolve to deploy more advanced and complex use cases or decrease the level of human involvement.

## HOW DOES GENAI COMPARE TO OTHER AUTOMATION TECHNOLOGIES?

Automation technologies, such as robotic process automation (RPA), have been used for several years by accounting and financial reporting professionals to automate routine and repetitive tasks. While automation technologies can be beneficial to automate tasks that are performed the same way every time, they typically cannot handle situations where the format or structure of data is different from how it was programmed. GenAI can address these limitations by providing the ability to accept unstructured inputs with greater variation. Since genAI has the potential to integrate with other technologies, task automation may look very different when using genAI compared to traditional automation using RPA that is focused on replicating repetitive tasks.

### CONSIDERATIONS FOR AUDITORS

Automation technologies and genAI have different risks; therefore, it is important that auditors understand the type of technology that a company is using. Auditors may ask:

- + Is the technology rules-based (i.e., performs the task the same way each time) or is it probabilistic (i.e., involves a degree of variation and is not programmed to perform the task the same way each time)?
- + Does the technology only accept inputs in a specific format, or can it accept unstructured inputs?
- + For processes that have been automated, is the company using RPA or similar automation technologies, genAI, or a combination?



# The Regulatory Environment

As the use of AI technologies, including genAI, evolves, there have been increased calls globally for stronger regulations related to the safe and responsible development and use of AI, including genAI. Although existing regulations in many countries already govern the use and protection of data or emerging technologies and are applicable to AI, many countries have also begun to adopt new regulations and frameworks specifically to mitigate security and safety risks of AI as well as to advance the ethical and responsible use of AI. The regulations and voluntary frameworks discussed herein are not all inclusive.

## WHITE HOUSE EXECUTIVE ORDER ON SAFE, SECURE, AND TRUSTWORTHY ARTIFICIAL INTELLIGENCE

In October 2023, President Biden issued an executive order (Executive Order 14410) focused on seizing opportunities presented by AI and managing the related risks.<sup>11</sup> Among other things, the executive order directs federal agencies to establish new standards for AI safety and security, protect data privacy, advance equity and civil rights, support workers, promote innovation and competition, and establish the government's own responsible AI program. It requires the developers of certain powerful AI systems to share safety test results and other critical information about those AI systems with the federal government and directs the establishment of rigorous standards for testing AI systems to ensure their safety before public release. In addition to existing state and federal laws that govern the use of data or AI technologies, the executive order is a significant step towards more robust AI regulation in the US.

## EU ARTIFICIAL INTELLIGENCE ACT (AI ACT)

In March 2024, Members of the European Parliament approved the EU AI Act, a comprehensive legal framework for regulation of the development and use of AI systems, including general purpose AI.<sup>12</sup> It employs a risk-based approach that prohibits certain uses of AI such as social scoring based on social behavior or personal characteristics and requires the development of appropriate guardrails to mitigate risks to society for high risk AI systems around data quality, transparency, testing, monitoring, reporting, security, human oversight, and accountability. This regulation will impact US companies that operate in the EU or develop AI models that are used in the EU.

## VOLUNTARY RISK MANAGEMENT FRAMEWORKS

In addition to complying with new regulations, some companies are applying the principles of voluntary AI risk management frameworks to responsibly use AI, including genAI. The National Institute of Standards and Technology (NIST) published the *AI Risk Management Framework*, which can be voluntarily used by organizations to incorporate

### CONSIDERATIONS FOR AUDITORS

New regulations may result in changes to a company's ICFR that could impact the audit. For example, new regulations may necessitate additional entity-level controls for the governance of AI as well as updated policies and procedures related to the training, development, use, and ongoing monitoring of AI technologies. Auditors are also responsible for considering certain laws and regulations and the possibility of illegal acts by companies.<sup>13</sup>

<sup>11</sup> FACT SHEET: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence | The White House

<sup>12</sup> Regulatory framework proposal on artificial intelligence | Shaping Europe's digital future (europa.eu)

<sup>13</sup> As required by PCAOB AS 2405, Illegal Acts by Clients.

trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems.<sup>14</sup> It provides a framework to manage risks arising from AI that could affect individuals, organizations, and broader society. The framework is structured around four key pillars – govern, map, measure, and manage – designed to help organizations identify and assess potential risks associated with AI. Further, in response to the executive order described above, NIST is expected to develop additional guidelines, standards, and processes for AI safety and security, including topics related to genAI risk management, AI evaluation, and security testing. Frameworks and guidelines in this area are rapidly evolving.

Additionally, COSO released the *Realize the Full Potential of AI: Applying the COSO Framework and Principles to Help Implement and Scale AI* guidance, which is designed to help companies apply the COSO ERM Framework to the use of AI.<sup>15</sup> Specifically, the guide focuses on the need for organizations to design and implement governance, risk management, and oversight strategies and structures to realize the potential of humans collaborating with AI. The International Organization for Standardization (ISO) also published several voluntary standards related to mitigating risks arising from AI, including ISO/IEC 23894, which provides guidance on AI-related risk management for organizations, and ISO/IEC 42001, which specifies requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System (AIMS).<sup>16</sup>

Finally, some companies are developing their own principles for ethical and responsible use of AI. These principles focus on key concepts such as accountability, reliability, transparency, explainability, and security, among others.

<sup>14</sup> AI Risk Management Framework | NIST

<sup>15</sup> Artificial Intelligence | COSO. The COSO ERM Framework differs from the COSO Internal Control – Integrated Framework. The ERM framework focuses on broader strategic objectives than the Internal Control – Integrated Framework.

<sup>16</sup> ISO/IEC 23894:2023 - Information technology – Artificial intelligence – Guidance on risk management and ISO/IEC 42001:2023 - Artificial intelligence – Management system

# Considerations When Auditing Companies That Are Deploying GenAI

## POTENTIAL RISKS ARISING FROM DEPLOYING GENAI

As the auditor obtains an understanding of how genAI is used in financial reporting processes and ICFR and the overall governance and oversight of genAI, the considerations described in the table below may help the auditor determine how the company's use of genAI technologies may impact the auditor's identification and assessment of risks of material misstatement, including the identification of process level risks or risks arising from IT. The considerations described below are not all-inclusive and will vary based on the company's facts and circumstances.<sup>17</sup>

Potential Risk Area	Example Risks or Sources of Risks	Questions for Auditor Consideration
<b>Governance</b>	+ AI solutions are not identified and managed appropriately and consistently across the company.	<ul style="list-style-type: none"> <li>+ Who (individual or group) in the company is responsible for oversight of the use of genAI?</li> <li>+ Has the company developed a framework for responsible use of genAI?</li> <li>+ Has the company established policies regarding the acceptable and ethical use of genAI?</li> <li>+ How are policies regarding acceptable and ethical use of genAI documented and communicated to appropriate individuals throughout the company?</li> <li>+ How does the company monitor compliance with policies regarding acceptable and ethical use of genAI?</li> <li>+ Does the company have a process to track and monitor the use of genAI throughout the company, including use by third-party service providers?</li> <li>+ How does the company evaluate the impact (nature and affected groups) of genAI technologies being deployed?</li> <li>+ How does the company track risks arising from the use of genAI technologies and mitigating responses?</li> </ul>

<sup>17</sup> The considerations described herein are not necessarily unique to genAI technologies and may also be applicable for other types of artificial intelligence.

Potential Risk Area	Example Risks or Sources of Risks	Questions for Auditor Consideration
<b>Regulatory</b>	<ul style="list-style-type: none"> <li>+ The company's use of genAI technologies violates contractual agreements, laws, or regulations.</li> </ul>	<ul style="list-style-type: none"> <li>+ What are the applicable laws and regulations impacting the company's use of genAI technologies?</li> <li>+ Do the company's policies and procedures to monitor compliance with laws and regulations include newly enacted and changes to existing laws and regulations related to genAI?</li> <li>+ Does the company have contractual agreements that may impact how the company can use genAI technologies?</li> <li>+ Has the company performed a regulatory, legal, and contractual compliance assessment to understand considerations for the design, deployment, and use of genAI technologies?</li> <li>+ If the company uses genAI technologies developed by a third party, is the company able to obtain sufficient information from the third-party provider regarding compliance with applicable laws, regulations, and contractual obligations?</li> <li>+ How does the company monitor genAI technologies over time to determine if bias has been introduced through the algorithms or the data that could result in noncompliance with laws, regulations, and contractual obligations?</li> </ul>
<b>Knowledge and Skills</b>	<ul style="list-style-type: none"> <li>+ Individuals in governance or management positions do not have the appropriate knowledge and skills to provide effective oversight of the company's approach to deploying genAI.</li> <li>+ The company does not have skilled resources to successfully oversee, develop, deploy, operate, and monitor genAI technologies.</li> <li>+ The company does not provide sufficient training to employees to use genAI technologies effectively and as designed or employees inappropriately rely on genAI technologies (automation bias).<sup>18</sup></li> </ul>	<ul style="list-style-type: none"> <li>+ Has the company identified specialized skills or knowledge needed to assist with oversight, development, deployment, operation, and monitoring of genAI technologies?</li> <li>+ How does the company provide training for employees and management who are responsible for oversight, developing, deploying, operating, or monitoring genAI technologies?</li> <li>+ How does the company educate employees and management on responsible use of genAI, including an understanding of the risks for AI hallucinations and guardrails on the ability to rely on the outputs?</li> <li>+ Does the company provide resources (such as user manuals and real-time support) for specific genAI technologies to employees?</li> <li>+ Does the company hire or engage third-party resources with the required expertise to help ensure successful oversight, development, deployment, operation, and monitoring of genAI technologies?</li> </ul>
<b>Fraud</b>	<ul style="list-style-type: none"> <li>+ GenAI technologies are used by employees, management, or third parties to perpetrate and conceal fraud.</li> </ul>	<ul style="list-style-type: none"> <li>+ How has the company considered genAI technologies in its fraud risk assessment?</li> <li>+ Has the company identified new incentives, opportunities, or pressures to commit fraud due to the deployment of genAI technologies?</li> </ul>

<sup>18</sup> Automation bias is a tendency to favor outputs generated from automated systems, even when human reasoning or contradictory information raises questions about whether such output is reliable or fit for purpose.

Potential Risk Area	Example Risks or Sources of Risks	Questions for Auditor Consideration
<b>Data Privacy</b>	<ul style="list-style-type: none"> <li>+ The company's confidential data is mismanaged because it is entered into a genAI technology (some third-party genAI technologies track and save all inputs to use for further development of the technology).</li> </ul>	<ul style="list-style-type: none"> <li>+ How does the company consider data privacy risks when selecting or developing genAI technologies?</li> <li>+ Does the company use a public instance of genAI technologies that tracks and saves inputs and data that are accessible by third parties or a private instance where inputs and data are tracked and saved only by the company?</li> </ul>
<b>Security</b>	<ul style="list-style-type: none"> <li>+ The company's genAI technology is susceptible to cyber-attacks, including data poisoning,<sup>19</sup> malicious prompt injections,<sup>20</sup> or malicious overriding of prompts.</li> </ul>	<ul style="list-style-type: none"> <li>+ How does the company consider cybersecurity risks when selecting or developing genAI technologies?</li> <li>+ Has the company performed a cybersecurity risk assessment to evaluate threats and safeguards?</li> </ul>
<b>Selection and Design of GenAI Technologies</b>	<ul style="list-style-type: none"> <li>+ The company selects or develops a genAI technology that does not achieve the desired objective.</li> </ul>	<ul style="list-style-type: none"> <li>+ How does the company identify appropriate processes that are suited for augmentation by genAI?</li> <li>+ How does the company design genAI technologies, including determining which genAI technologies to use (such as, selecting an existing genAI technology, using a foundation model with added customizations, or developing the company's own model) and the data needed for those technologies?</li> <li>+ How does the company select third-party genAI technologies for use?</li> <li>+ Has the company developed clear objectives and related success criteria for genAI technologies?</li> <li>+ How are genAI technologies configured within the company's IT environment?</li> </ul>
<b>Use of a Foundation Model</b>  <i>For genAI technologies that use a foundation model (with or without customizations from the company)</i>	<ul style="list-style-type: none"> <li>+ The foundation model is unreliable resulting in repeated errors or a favoring of certain results or outputs by the model.</li> </ul>	<ul style="list-style-type: none"> <li>+ How does the company consider whether the foundation model is appropriately suited for the company's needs?</li> <li>+ How does the company evaluate the model for bias?</li> <li>+ How does the company determine whether to add customizations to the foundation model to meet the company's specific needs?</li> </ul>

<sup>19</sup> Data poisoning involves deliberately providing genAI technologies with unreliable data to influence the initial training, ongoing learning, or future retrieval, leading the technology to provide unreliable outputs.

<sup>20</sup> Malicious prompt injections involve prompting genAI technologies to provide unreliable outputs. Malicious prompt injections can be direct (a user provides a malicious prompt to the genAI technology) or indirect (malicious prompts are hidden in or disguised as data).



Potential Risk Area	Example Risks or Sources of Risks	Questions for Auditor Consideration
<p><b>Model Training and Development</b></p> <p><i>Applicable for genAI technologies that use a model developed by the company and for incremental customizations to a foundation model by the company</i></p>	<ul style="list-style-type: none"> <li>+ The methods used to train the genAI model are insufficient or otherwise not appropriate resulting in repeated errors or a favoring of certain results or outputs by the model.</li> <li>+ The training of the genAI model introduces biases of the human programmer resulting in repeated errors or a favoring of certain results or outputs by the model.</li> <li>+ The data used by the company to train the model is biased or otherwise not reliable resulting in repeated errors or a favoring of certain results or outputs by the model.</li> </ul>	<ul style="list-style-type: none"> <li>+ How does the company evaluate the sufficiency of training of the genAI model?</li> <li>+ How does the company evaluate the model for bias?</li> <li>+ How does the company evaluate the training data for reliability and data quality?</li> </ul>
<p><b>Model Performance</b></p>	<ul style="list-style-type: none"> <li>+ GenAI technologies do not consistently operate in accordance with their intended purpose and at an appropriate level of precision.</li> <li>+ GenAI technologies provide incomplete, inaccurate, or unreliable outputs (AI hallucinations).</li> <li>+ GenAI technologies provide outdated or other information that is not relevant.<sup>21</sup></li> </ul>	<ul style="list-style-type: none"> <li>+ How does the company test genAI technologies prior to deployment to determine that they operate as designed?</li> <li>+ How does the company assess the relevance and reliability of genAI outputs for the intended purpose?</li> <li>+ Does the company measure, track, and communicate performance metrics related to the functioning of the genAI technologies, including the precision of the technology?</li> </ul>
<p><b>Prompts</b></p>	<ul style="list-style-type: none"> <li>+ Prompts entered into genAI technology by employees are not appropriate to achieve the intended output from the genAI technology.</li> </ul>	<ul style="list-style-type: none"> <li>+ How has the company trained employees operating genAI technologies about appropriate prompts?</li> <li>+ Does the company have standardized prompts for employees to use when operating genAI technologies?</li> <li>+ If prompts include data, how has the company considered the reliability of data used in the prompt?</li> </ul>
<p><b>Ongoing Reliability and Monitoring</b></p>	<ul style="list-style-type: none"> <li>+ GenAI technologies are not monitored after deployment to determine whether they are functioning appropriately.</li> <li>+ After deployment, genAI technologies do not continue to function as designed due to the technologies' evolution over time or to intentional or unintentional changes to genAI technologies.</li> </ul>	<ul style="list-style-type: none"> <li>+ How does the company monitor the ongoing effectiveness of genAI technologies for the intended purpose?</li> <li>+ Does the company have a process to periodically reevaluate genAI technologies to determine whether they are functioning as intended?</li> <li>+ How does the company monitor changes to genAI technologies?</li> </ul>

<sup>21</sup> GenAI technologies often do not have access to real time data and information (as the data that the model is trained on is only through a specific point in time), and therefore, genAI technologies may state information that is correct based on its training data but is not currently relevant.

## RESPONDING TO IDENTIFIED RISKS

### Human in the Loop

For many current genAI use cases in financial reporting processes and ICFR, keeping a human involved in the process (“a human in the loop”) may address some of the risks arising from its use. Keeping a human in the loop means that employees are responsible for performing the following, as appropriate, (a) reviewing the accuracy and completeness of company inputs entered into the genAI technology, (b) understanding the explainability and interpretability of the outputs from the genAI technology, and (c) reviewing the outputs from the genAI technology to determine their quality, reliability, and appropriateness. Generally, keeping a human in the loop can support the identification of inaccuracies, including incomplete output from the genAI technology. The level of human involvement, including the review of inputs and outputs, may vary depending on the genAI use case, is commensurate with the risk profile and environment that the genAI technology operates in, and may evolve over time. Human involvement with genAI technology requires a high degree of vigilance and professional skepticism.

### Audit Response: Internal Control Considerations

Based on the auditor’s risk assessment, the auditor may determine whether it is necessary to test certain control activities related to the company’s use of genAI.<sup>22</sup> When there is a human in the loop, an auditor’s evaluation of the design and operating effectiveness of such control activities may include the sufficiency and appropriateness of the employee’s review of the completeness, accuracy, and relevancy of the output from the genAI technology. For example, auditors may consider if the reviewer appropriately considered the unique risks related to genAI when performing their review. It is important to note that control activities related to the human involvement with genAI are supported by appropriate entity-level controls or general IT controls.

### CONSIDERATIONS FOR AUDITORS

Auditors may consider the following questions related to human in the loop involvement:

- + How does the company determine the appropriate level of human in the loop involvement with genAI technologies?
- + How does the company develop processes to promote appropriate human in the loop involvement in reviewing outputs from genAI technologies?
- + How does the company consider explainability and interpretability needs of users to enable effective human in the loop involvement with the genAI technology?

<sup>22</sup> When performing an integrated audit in accordance with PCAOB AS 2201.39, “[t]he auditor should test those controls that are important to the auditor’s conclusion about whether the company’s controls sufficiently address the assessed risk of misstatement to each relevant assertion.”

# Example Use Cases



While use cases will vary based on a company's operations, processes, and specific facts and circumstances, the following examples demonstrate how auditors may encounter genAI in a company's financial reporting processes and ICFR.

## Drafting Financial Statement Disclosures

Company X includes required property, plant, and equipment (PP&E) disclosures in its annual financial statements. Previously, an employee involved in the company's financial reporting process prepared the draft disclosure based on underlying supporting schedules and general ledger data. To enhance the efficiency of the process, Company X deployed genAI technology to prepare the first draft of the disclosure using the prior year disclosure, underlying schedules, and data from the general ledger. The financial reporting employee is now responsible for reviewing the draft disclosure prepared by the genAI technology. The disclosure then goes through the existing review process in which the assigned reviewer considers that the work was prepared by genAI and the associated risks. While genAI is used in the process, there is still a high level of human involvement as the employee reviews and verifies the disclosure drafted by the genAI technology.

In its risk assessment, Company X may identify risks related to model performance, among other risks. For example, as it relates to model performance, Company X identified two risks:

1. The disclosure prepared by the genAI technology is not complete or accurate based on the underlying data.
2. The disclosure prepared by the genAI technology is not appropriate because there have been updates to US GAAP PP&E disclosure requirements after the cut-off of the model's training data.

Company X determined that both risks are mitigated by the following existing control, with certain enhancements to address that the work was performed using genAI:

1. Company X has a control in which an employee with appropriate authority and competence (i.e., knowledge of US GAAP and the disclosure requirements related to PP&E) reviews the draft PP&E disclosure to validate that the disclosure is complete and accurate, agrees to the underlying schedules, and includes all information required by US GAAP.

Note that there may be additional ICFR considerations addressed through entity-level controls or general IT controls.

While genAI is used in the process, there is still a high level of human involvement as the employee reviews and verifies the disclosure drafted by the genAI technology.

## Drafting Code for Reports

Company Y has an internal control (control A) whereby all significantly aged receivables are reviewed by an individual with appropriate authority and competence to evaluate their collectability and assess the appropriateness of the related allowance, or lack thereof. In the performance of control A, the control operator relies on report A, which lists key attributes for all outstanding receivables aged greater than 30 days, including customer name, receivable amount, days outstanding, and allowance amount (if applicable), among other attributes.

The control operator generates report A using a report writer tool based on SQL code. Previously, if the control operator needed modifications to the report, an employee would write the code. Now, when updates are needed, an employee uses genAI to write the SQL code. The employee, who has appropriate expertise in SQL, reviews the code drafted by the genAI technology before the updates are made. Humans remain in the loop to review and verify the appropriateness of the code drafted by genAI, and the code goes through the normal testing protocols prior to those changes being finalized.

With respect to the SQL code, Company Y may identify risks arising from the use of genAI, such as model performance. Specifically, Company Y identified a risk that the code prepared by genAI does not produce complete and accurate results. Company Y determined that this risk is mitigated by the following existing controls:

1. Company Y has a control in which an employee with appropriate authority and competence (i.e., knowledge of SQL code) reviews and approves any changes to report codes prior to implementing the changes.
2. Company Y has a control in which the report code is tested in a non-production environment prior to the code being implemented into production.

Note that there may be additional ICFR considerations addressed through entity-level controls or general IT controls.

### FUTURE STATE

In the future, as Company Y increases reliance on genAI, there is a potential for this example to evolve to the point where the control operator prompts a genAI technology to make desired changes to the report (e.g., additional fields added) and the genAI technology prepares the code and, with additional programming or interfaces, puts it into production without any human involvement.

# Additional Audit Considerations



## **KNOWLEDGE AND SKILL OF THE AUDIT ENGAGEMENT TEAM**

When auditing companies that are deploying genAI technologies, auditors will consider whether the audit engagement team has the appropriate knowledge and skills to identify, evaluate, and respond to risks of material misstatement, including process level risks or risks arising from IT, related to the company's use of genAI. Based on the skillset of the audit engagement team members, it may be necessary to complete additional training related to genAI technologies. In other cases, the audit engagement team may determine that it is appropriate to involve an individual with the requisite knowledge, skill, and ability related to genAI.

## **FUTURE STATE**

As the use of genAI technologies evolves and companies place more reliance on genAI technologies, audit procedures will likely need to evolve as well. For example, the nature of outputs from genAI technology may be complex or otherwise unable to be independently replicated or verified by a human in the loop or tested by an auditor. In these cases, it will be critical for companies to have appropriate processes and controls surrounding the use of genAI, including human oversight of the genAI technology, which are designed to respond to the associated risks. To obtain sufficient appropriate audit evidence, tests of controls are required when substantive procedures alone are not sufficient.<sup>23</sup> Auditors will need to consider the implications of the company's use of genAI technologies on tests of controls and substantive procedures in making a conclusion about whether sufficient appropriate audit evidence has been obtained. Obtaining sufficient appropriate audit evidence when companies place increased reliance on genAI technologies will be an area of continued focus and potential challenge for auditors.

# Conclusion



The use of genAI in financial reporting processes or ICFR by companies introduces new risk considerations for auditors. It is important for auditors to be mindful of the risks and challenges that can arise from a company using genAI. Auditors are well-suited to apply and build on their expertise in identifying and assessing risks, exercising professional skepticism, and developing appropriate audit responses.

<sup>23</sup> Refer to PCAOB AS 2301.17.



# CAQ

[www.thecaq.org](http://www.thecaq.org)

**We welcome  
your feedback!**

Please send your comments or  
questions to [hello@thecaq.org](mailto:hello@thecaq.org)