Deloitte.

Navigating the artificial intelligence frontier

An introduction for internal audit

Insights to ground your AI strategy

During 2023, artificial intelligence (AI) captured the imagination of the world, fueling discussion among businesses and policymakers by demonstrating the transformative power of how these technologies could redefine work. Whether interacting with vast sources of knowledge and business data through humanlike interactions, accelerating how people work, or revealing new opportunities that were not previously possible through manual efforts, the benefits of AI are far reaching.

AI is a broad domain. However, significant attention has been given to a specific field of AI known as Generative Artificial Intelligence (GenAI) following the mass global interest in applications like OpenAI's ChatGPT. Adoption and use of (GenAI) has been sudden and rapid among the public. OpenAI reported reaching 100 million users within 60 days of releasing ChatGPT to the public.[1] Given the opportunity GenAI presents, and the fact employees are using GenAI "side of desk" for work tasks, it is no wonder that organizations are investing heavily in enterprise use cases. With the rapid acceleration and integration of GenAI into business functions, AI and GenAI risk management will continue to be a hot topic for internal audit teams throughout 2024 and beyond.

For internal audit, this presents two key considerations: how to provide assurance and assess the risks associated with AI (including GenAI), and how to leverage its potential to evolve and innovate internal audit's own ways of working. In this publication, we explore these two faces of AI.

Internal audit's role in assurance over AI

AI and GenAI offer significant opportunities for organizations. At the same time, they present a frontier of new risks for boards and audit committees to navigate. To mitigate and minimize these risks, organizations are actively investing in the development of risk management frameworks and controls to enable them to innovate with confidence.

These new AI controls will be needed to help manage data privacy and security risks, as well as ethical considerations and concerns about the reliability of outputs created by GenAI. Internal audit functions are also looking at the developing regulatory landscape to ensure that their organization is preparing for the arrival of regulations across all geographies they operate in.

In conjunction with the publication of the regulations and guidance, the pace of AI development and deployment for the United States is expected to intensify as the US government pushes to be a global leader in AI development and innovation.

Harnessing the power of AI to reimagine internal audit's ways of working

Alongside their organizations' efforts to leverage AI, internal audit leaders are also trying to understand the potential impact or opportunity and art of the possible for their own functions.

As a firm, we believe the integration and use of enabling technologies in internal audit, such as AI, is critical to helping functions maximize their impact and value. The digital landscape is broad, covering many other domains including automation, audit management systems, cloud-based solutions, visualization, data analytics, and process mining. While they can be deployed in isolation, the power of digital is in their combination. As such, internal audit functions need a strategic and coordinated approach across both the function and the internal audit life cycle.

GenAI will play an important role within internal audit functions' digital strategies, not only in providing new capabilities but also in helping to engage leadership and staff in continuous improvement and innovation by reimagining traditional approaches. For those who are successful in digitizing their functions, the rewards are clear—enhanced quality, increased assurance and better insights, new levels of productivity, increased staff satisfaction levels, and greater impact on the organization and for their broader stakeholders. Yet digital capabilities remain a significant gap and the number-one opportunity for many functions.

[1] Krystal Hu, "ChatGPT sets record for fastest-growing user base – analyst note," Reuters, February 2, 2023.

Decoding the jargon: Useful AI terminology to know

Before internal audit functions can hope to assess and assure the risks of AI or look to explore the art of the possible from its use, internal auditors must acquire a baseline of AI fluency. To help you with your AI 101, we outline some of the key terminologies and basics below.

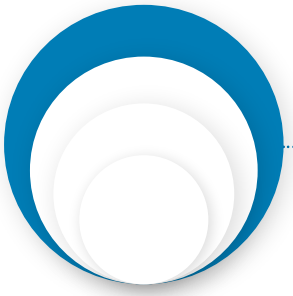

Artificial intelligence

Artificial intelligence is a broad "umbrella" term given to the field of computer science that focuses on creating systems that can perform tasks requiring human intelligence.

Machine learning

Machine learning refers to algorithms that make informed decisions and learn over time without being explicitly programmed to do so. Machine learning helps to train AI models to identify and predict patterns based on human-processed data, rather than relying on hard-coded rules.

Deep learning

Deep learning is a powerful and advanced machine learning paradigm that leverages neural networks to improve model performance. The models simulate human reasoning to make intelligent decisions and learn over time based on observed results.

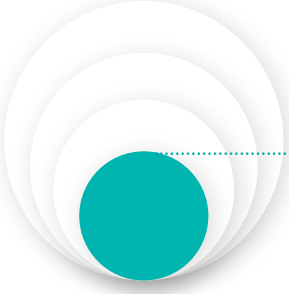

Generative artificial intelligence

GenAI is a highly sophisticated subset of AI using foundation models to create content across a variety of modalities. The models which support the generation of content (often referred to as foundation models) are underpinned by advanced machine and deep learning capabilities.

Types of GenAI (modalities)

The primary GenAI foundation models are focused around generating new content using our primary forms of communication, such as text and imagery. However, there are many variations, and models continue to develop at pace. For simplicity, it can be helpful to think of the models as being able to produce outputs across an increasing combination of the following data modalities (e.g., multimodal):

Explain to my colleagues the business impact of Generative AI in 50 words

Text

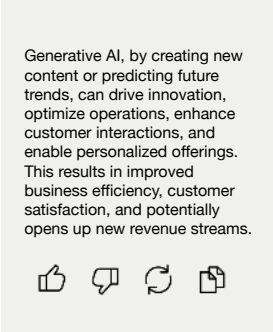

In Python, code a program that predicts the likelihood of customer conversion

Code

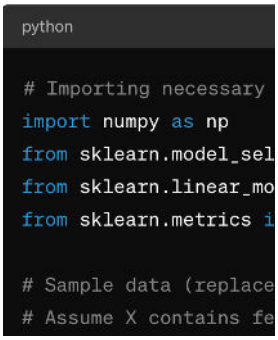

A bowl of soup that is a portal to another dimension as digital art

Image

A teddy bear painting a portrait

Video

Play "we have to reduce the number of plastic bags" in a sleepy tone

Audio

Large language models

Of all the modalities, large language models (LLMs) have gained most of the attention from organizations for their ability to generate text. Popularized through tools like OpenAI's ChatGPT, LLMs are a specific type of text-based model that have been trained on petabytes worth of global data. The parameters within these models represent the model's level of understanding about each word and their context within the training datasets. In the case of LLMs, more parameters allow them to capture more complex patterns in the data they were trained on, typically leading to improved accuracy on language-related tasks.

At a simplistic level, LLMs predict an output based on inferences built on their training and the inputs they receive. Given the level of data they have been trained on, their ability to provide highly convincing and compelling responses in a humanlike interaction is why they have gained such attention.

The mechanics of GenAI

GenAI is a highly sophisticated subset of AI. While the vast majority of internal auditors will not need a deep technical understanding akin to data scientists, it can be helpful to appreciate the general mechanics of GenAI to consider where risks can arise and to determine the level of technical skills that an internal auditor may need to provide assurance over the organizations' use of these tools.

How GenAI works

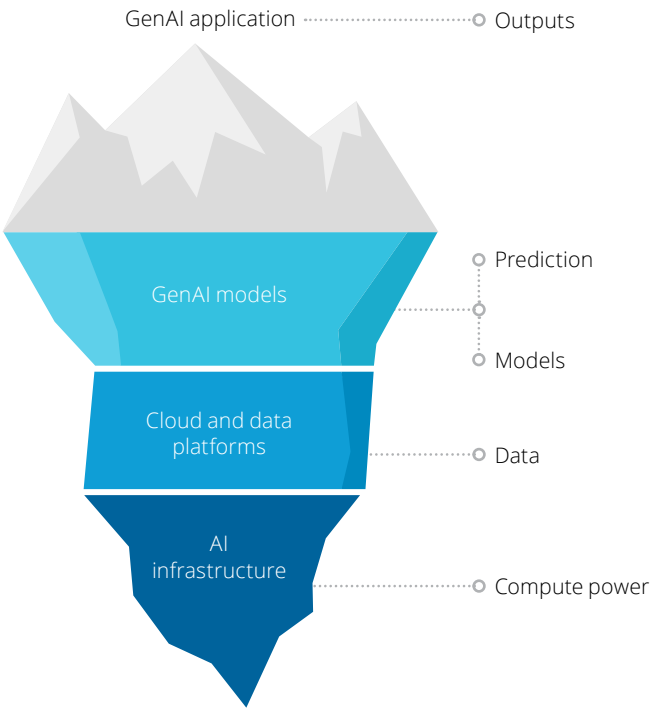

Applications ... what we see

GenAI applications generate content from user prompts across various modalities (e.g., text, image, video, audio) based on how the underlying model was trained.

Why do these applications seem so human?

Like traditional AI, foundation models are models (1) that predict outputs based on inferences built on the inputs they receive. However, through fine-tuning (2), prompt engineering (3), and adversarial training (4), these models can generate outputs geared toward meeting human intent.

What are foundation models?

OpenAI's GPT-4 and NVIDIA's Megatron are two examples of foundation models, specifically large language models, which use deep learning to process massive amounts of data to form "memories" on the input datasets through tokenization (5), thereby shaping the models' parameters (6). There are common foundation model architectures—for example, Transformer (7), Diffusion (8)—which drive the modalities for each model.

Training on the world's knowledge

Foundation models are trained on petabytes worth of global data to shape understanding, tone, and behavior while considering human communication styles.

Powering our journey to tomorrow

The scale of the compute capacity required to train and process foundation models necessitates the usage of leading GPUs (graphics processing units) (e.g., A100 NVIDIA) and TPUs (tensor processing units) (e.g., Google TPU v4) on scalable infrastructure.

Key terms

1. **GenAI model**
   A neural network that has undergone training to generate outputs based on a given input prompt.

2. **Fine-tuning**
   The process of refining foundation models to make them suitable for specific applications.

3. **Prompt engineering**
   The act of creating or modifying the prompt given to a model to obtain an optimal answer or output.

4. **Adversarial training**
   The technique of pitting different deep learning models against each other in a training game or competition.

5. **Tokenization**
   The process of splitting text into smaller units.

6. **Parameters**
   Trainable values within the model that are adjusted based on the training data to optimize the output.

7. **Transformer model**
   A model that can "transform" words into context-aware representations that Google and University of Toronto invented in 2017.

8. **Diffusion model**
   Construction of high-resolution images from noise. Mostly used in speech-to-image and text-to-image models.

What can AI do today?

The capabilities that AI can provide today are allowing organizations to challenge their ways of working and reveal new possibilities. Not all of these will be relevant to internal audit, but some could be applied across the internal audit life cycle to evolve and innovate approaches. Example AI capabilities include:

SENSE	
Sense physical data	**Sense visual data**
Sense light	Sense screen pixels
Sense sound	Sense keystrokes
Sense temperature	Sense mouse clicks

PERCEIVE			
See objects	**See faces**	**See actions**	**Hear voices**
Detect objects	Detect faces	Detect motion	Convert speech to text
Classify objects	Recognize faces	Identify actions	Identify speaker
Perform OCR	Determine age	**Hear sounds**	Determine gender from voice
	Determine gender	Recognize sounds	Identify emotion in voice
	Recognize emotion		

LEARN	
Learn by technique	**Learn facts and skills**
Learn from examples	Learn skills
Learn by trial and error	Learn facts
Learn by analyzing structure	

KNOW	
Represent and store knowledge	**Retrieve information**
Populate global knowledge base	Retrieve relevant documents
Populate contextual knowledge base	Retrieve relevant answer units
Maintain truth	Retrieve specific facts

Capabilities with most relevance and potential application to internal audit.

COMMUNICATE			
Understand language		**Understand and generate language**	
Classify text	Detect language	Translate languages	Generate narrative
Extract entities	Analyze sentiment in text	Answer questions	Generate image and video captions
Recognize relationships	Analyze emotion in text	Dialogue	

PLAN	ACT
Plan production	**Act in physical environment**
Plan robot motion	Convert text to speech
Plan routes	Move robot limbs
	Act in virtual environment
	Generate mouse clicks & keystrokes
	Generate animated avatar

CREATE	
Create text	**Create videos**
Create marketing content	Create custom videos
Create sales content	
Create support content	
Create images	**Create speech**
Create general images	Create custom voices
Create advertising images	
Create models	**Create chemicals**

REASON AND SOLVE PROBLEMS	
Infer	
Make logical inferences	Cluster
Make probabilistic inferences	Recommend
Classify	Predict numeric value
Solve problems	
Search for optimal solution	Optimize
Satisfy constraints	

Capabilities with most relevance and potential application to internal audit.

AI is not as new as you think … but with GenAI we are heading into unchartered waters

Before becoming too caught up in the GenAI hype, it is worth recognizing that most people are already using AI in their daily lives without realizing it. For example, tools like auto-complete, spellcheck, smart calendar scheduling, and suggestions on the most effective ways to visualize data in applications, such as Power BI, are all using forms of AI. Natural language processing (e.g., chatbots, sentiment analysis), speech recognition (audio to text), robotics, and perception sensing (e.g., object detection) have been in existence for some time. Chances are your organization, and potentially your own internal audit function, are already engaged in forms of machine learning. If you have not yet explored existing AI capabilities, there are significant opportunities and benefits that can be gained before heading into the world of GenAI. The clear potential from democratizing access and the acceleration in development of GenAI tools is creating very significant opportunities and brings with it new areas of risk that many organizations have yet to understand.

Internal audit's role in assurance over AI

GenAI presents a broad spectrum of risks, many of which are still emerging. Among the main concerns raised by GenAI are:

Risk	Description
Privacy	Personal information shared with third-party Software-as-a-Service AI may not comply with privacy laws and puts customer/employee data at risk of exposure.
Intellectual property	Information gathered (e.g., by web scraping) may contain IP protected content, and prompts must be carefully written not to leak any secret know-how. There are also challenges with protecting IP of content generated by AI.
Malicious behavior	GenAI tools may be targeted by adversaries to reveal sensitive information and/or take malicious actions on networks and data.
Ethical use	GenAI tools may be used in an unintended manner and to circumvent organizational policies, laws, and regulations (e.g., submitting content in competitive events).
Hallucination	Models might output facts that are false. Sources and citations are unavailable for most models.
Bias	Bias in training data (e.g., over/under-representation of a population cohort, sexism, racism) can generate biased outputs.
Model performance	Lack of considerations for model performance limitations (dependent on training data used) could lead to sub-optimal business outcomes (e.g., poor quality reports).

The regulatory landscape

There have been several developments in the AI regulatory landscape, which continues to move at pace. Guidance has been published to aid organizations as they navigate the use not only of GenAI, but of all forms of AI. Some of the key voices in the regulatory landscape include:

- **EU AI Act (latest development from December 2023)** – The European Union AI Act, which was provisionally agreed to among member states and is expected to come into action in the first quarter of 2024, is a regulatory risk-based approach to classify AI systems and manage the development, distribution, and use of AI systems.
- **US White House Executive Order on AI (issued October 2023)** – The Biden-Harris administration has released an executive order (EO) aimed at enhancing safe, secure, and trustworthy development and use of AI throughout the federal government. This EO, while not a law or regulation, encourages federal agencies to explore AI uses responsibly and manage associated risks, and could lead to new policies impacting AI developers. To advance security and safety in AI's development and use, the EO invokes emergency authority to require disclosure of powerful AI systems and large-scale computing operations. It also addresses concerns around GenAI, encouraging the identification of synthetic content and the use of labels to distinguish between authentic and AI-generated content.
- **International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC) published two key standards for managing AI risks and systems (published in February 2023 and December 2023)** – ISO/IEC 23894:2023 (published in February 2023) provides risk management guidance for organizations developing, deploying, or using AI systems. It outlines principles and processes for integrating risk management practices throughout the AI life cycle. ISO/IEC 42001:2023 (published in December 2023) focuses on establishing, implementing, maintaining, and continually improving an AI management system. It specifies requirements to facilitate the responsible development, deployment, and use of AI.
- **NIST (National Institute of Standards and Technology) framework (published in January 2023)**[2] – NIST has collaborated with organizations from both public and private sectors to develop the NIST AI risk management framework. The guidance is voluntary and aims to help organizations understand the considerations that should be made during the design, development, use, and evaluation of AI systems.

[2] Marilena Do Rosario, "What you need to know about NIST's AI Risk Management Framework published in January 2023," Deloitte Financial Services Blog, February 2, 2023.

What should internal audit be doing for their boards and audit committees?

While GenAI technology is still developing, it is already being adopted by organizations at pace. Internal audit functions are seeking to understand to what extent their organization is using this technology and to what extent they are planning to invest in it. As internal audit functions grapple with this new risk domain, we recommend the following activities:

01	**AI strategy and governance** Internal audit should consider its organization's approach to the governance of AI. This should include a review of the organization's AI strategy, business case(s), and to what extent AI risks have been considered. Consideration should be given to what extent senior executives have been involved in defining the AI strategy and associated guardrails, given they can have organizational consequences.
02	**Policy, standards, and guidelines** Internal audit should consider reviewing any AI policy the organization has developed, including acceptable usage guidance and/or policy that defines the parameters of AI system development.
03	**AI inventory** Internal audit should consider whether an AI inventory has been developed by the business including both active and developing AI projects with details on their status and risk management considerations. Organizations are taking differing approaches to this, but ultimately AI risks cannot be managed unless there is clarity over AI use.
04	**Regulatory readiness** Internal audit should understand how the organization is staying up to date with the fast-moving regulatory environment. Organizations need to consider regulations in all the geographies they operate in. If this assessment is not thorough, they run the risk of having to "roll back" deployed AI use cases, which could cause significant business disruption.
05	**AI risk management and culture** Current risk management processes may need to be amended to ensure that risks associated with AI are proficiently covered.[3] AI risk management frameworks and risk assessments are being developed and should be integrated into the current risk management processes and procedures to ensure systems utilizing AI are effectively managed, governed, and monitored. Risk appetite statements may also need to be updated, and many organizations are adapting existing governance arrangements to be fit for AI, such as AI ethics councils and the creation of AI centers of excellence.

[3] Lukas Kruger and Michelle Seng Ah Lee, "Embedding controls and risk mitigations throughout the GenAI development lifecycle," Deloitte, accessed March 28, 2024.

Harnessing the power of AI to reimagine internal audit's ways of working

What about internal audit's own use of AI?

The use of more established AI capabilities (e.g., natural language processing and machine learning) have been present within more advanced internal audit functions for some time, often found within analytics teams. As access to GenAI and data security issues are overcome, we expect to see internal audit functions of all shapes and sizes to significantly scale their use of GenAI. For now, the reality is that most internal audit functions have not engaged in GenAI beyond exploration of ChatGPT or conceptual applications. Only very few are actively developing proofs of concept, but this is just a matter of time. As AI continued to grow in popularity, it is hard to imagine a world where it is not at the forefront of our businesses. The good news is that AI is not as scary as it seems.

Enabling technologies are becoming increasingly accessible and this is only being accelerated through the wider efforts of organizational Information Technology (IT) functions looking at the same challenging questions. Organizations do not need to become digital experts overnight or start replacing auditors with a team of data scientists (although increasing digital fluency and being able to access some of these skill sets will be important).

A glimpse into the GenAI-driven internal audit life cycle

The application of GenAI on internal audit's life cycle is only limited by the imagination and creativity of teams. From our discussions with internal audit functions, the following applications and use cases are where many in the industry see potential:

Risk assessment	Plan development	Engagement planning	Execution	Reporting
Supporting auditor research and understanding of risk for a specific industry. Supporting audit universe creation, e.g., guidance on universe design and process universe.	Supporting auditor research and understanding of risk, business processes, and expected controls in advance of engagement planning. Suggested audits against the risk- assessed audit universe. Suggested scheduling and resource allocation based on known constraints, e.g., number of staff, their skills and seniority.	Supporting auditor research and understanding of risk and business processes in advance of planning. Suggested control objectives and test procedures based on in-scope risk areas. Suggested data sources and potential analytics tests. Generated scripts for data extraction and analytics execution. First draft of scope/terms of reference.	Analysis of data through natural language questioning. Suggested interview questions for different stakeholders' personas. Critical assessment of risk and control descriptions, e.g., if it covers who, what, where, and when. Initial draft of workpaper. Drawing themes from interview notes/audio. Summation/ interrogation of audit evidence documents. Initial workpaper review and quality assurance. Initial draft of issue/ observations.	Initial draft report. Initial draft report review and QA. Editorial QA, e.g., simplifying language, sentiment analysis. Summation of reports for audit committee summaries. Generation of video/ audio reporting. Customized stakeholder communications. Report language translation. Drafting emails to communicate the audit report.

AI is only one element of internal audit's digital landscape. Significant benefits can be achieved through automation, audit management systems configuration and design, cloud-based solutions, visualization, data analytics, and process mining. While they can be deployed in isolation, the power of digital is in their combination. As such, internal audit functions need a clear digital strategy and coordinated approach across the function. For further information on how functions should approach a purpose-driven and digitally powered future, we recommend reading our Internal Audit 4.0 framework.

What should internal audit be doing to accelerate its adoption of AI?

01	**Increase your digital fluency** Start engaging with learning and development now. You do not need everyone to become data scientists, engineers, or digital experts. However, being familiar with the terminology, types of capability, and potential for these tools will help accelerate adoption.
02	**Determine your digital strategy and potential** Determine how GenAI can help you achieve your broader functional strategy and outcomes. Systematically review your ways of working to identify potential use cases. But do not limit your digital strategy to just GenAI; there are many applications and use cases relating to other areas of machine learning, such as natural language processing, sentiment analysis, topic modeling, linear regression, and neural networks that can already be harnessed and provide opportunities for experimentation. Equally, do not overlook the opportunities that exist from maximizing functionality from audit management systems and embracing analytics, visualization, and other tools such as process mining.
03	**Engage with your technology teams** Understand your organization's stance toward AI, both from a data privacy and security perspective and its appetite for shaping existing solutions within the safety of your organization's environment.
04	**Clean up your data** The quality of AI both in terms of its training and its output will be a product of the quality of data it is given. Many organizations (including internal audit) have poor data quality, version control, or outdated versions of documents that have not been removed from intranets for years. While you are waiting for some of the tools to become more accessible, getting your house in order will pay dividends to the value AI can deliver. For example, analyzing your risk and control frameworks, scope documents, findings, and recommendations to create a tokenized database of internal audit content could help turn currently untapped information into a goldmine of knowledge and insight.
05	**Work through and manage the risks** The risks outlined in this publication are as relevant to internal audit's use of GenAI as they are to the business. Good governance is critical, and functions should be challenging themselves to put in robust governance processes and controls around the use, development, testing, access, and ongoing monitoring of AI within internal audit.
06	**Develop a culture of innovation** Organizational culture can make or break the success of innovative technology and new ways of working. The limits of what GenAI could be used for are only contained by the imagination of individuals. Functions that have a culture of innovation, curiosity, and the willingness to experiment have usually fared better than those that were less willing to embrace change. Functions should consider innovation programs, encourage experimentation, and reward the right behaviors.

Where do we go from here?

Whether it is the assurance over risks posed by AI or exploring how you might use these technologies in your own internal audit function, learning and improving your digital fluency is key. GenAI will also require a mindset shift. Its prevalence and the speed at which it is evolving will drive a need to reimagine the human-technology relationship. Interacting with GenAI will become part of the daily routine, enabling new possibilities but bringing the potential for overreliance on AI outputs.

Organizations, including internal audit functions, will need to assess the risks and opportunities associated with GenAI, balancing the benefits from efficiencies gained through reduced manual effort with the need to check and verify accuracy. What is clear is that tools this powerful offer so much potential that they will be here for the long term. The attention given to GenAI and the investments being made means internal audit will need to engage and do so quickly.

GenAI is a fast-moving and developing field of AI. As an organization, we are taking the same journey as many of our clients. We believe GenAI has the potential to transform the internal audit profession and have already made significant investments in both our approach to assuring GenAI and supporting organizations in their use of these technologies. If you would like to talk to our dedicated team of specialists, please get in touch.

Internal audit contacts

Sarah Fedele
Principal
US Internal Audit Leader
Deloitte & Touche LLP
sarahfedele@deloitte.com

Neil White
Principal
Deloitte & Touche LLP
nwhite@deloitte.com

Michael Schor
Principal
Deloitte & Touche LLP
mschor@deloitte.com

Geoffrey Kovesdy
Principal
Deloitte & Touche LLP
gkovesdy@deloitte.com

GenAI internal audit contacts

Alex Vorpahl
Senior Manager
Deloitte & Touche LLP
avorpahl@deloitte.com

Madeline Mitchell
Senior Manager
Deloitte & Touche LLP
mademitchell@deloitte.com

Explore Internal Audit 4.0 framework for further insights on the power of digital for internal audit

Deloitte.

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.

As used in this document, "Deloitte" means Deloitte & Touche LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Copyright © 2024 Deloitte Development LLC. All rights reserved.