

**Deloitte.**



## **Embracing GenAI**

A Comprehensive  
Compliance Approach  
for Generative AI Systems

# Introduction

Artificial Intelligence (AI) is increasingly becoming interwoven into the operational fabric of numerous organizations. Two dominant forms, Predictive AI and Generative AI (GenAI), offer different capabilities but also pose distinct challenges, particularly in the realm of compliance. This whitepaper aims to delineate the differences between Predictive AI and GenAI, the compliance considerations for GenAI, and how tools and strategies such as Deloitte's Trustworthy AI™ framework coupled with services from Amazon Web Services (AWS) such as Amazon Bedrock, AWS Audit Manager, and Deloitte Nexus Digital Nerve Center™ ("Deloitte Nexus") can provide leading practices to help organizations navigate the intricate compliance landscape of GenAI. The whitepaper delves into the amplification of data issues, model security concerns, model drift issues, the need for contextual measures to prevent bias, and how a broad approach to managing GenAI compliance can empower organizations to harness the transformative power of GenAI effectively and efficiently.



# Predictive AI vs Generative AI

AI is a complex and rapidly evolving field. The emergence of GenAI into the public sphere gained consciousness with language-learning models in 2020 and has continued to evolve. The increase of individual use of AI models represented a significant shift in public awareness and engagement with AI, catalyzing the integration of AI technologies into expansion of new use cases and everyday activities. AI has evolved significantly from Predictive AI, which was primarily rule-based and focused on specific tasks such as voice recognition, image recognition, or natural language processing. This form of AI, often based on expert systems and machine learning, relied on structured data and predefined algorithms to make decisions. Due to the development of transformers, large language models were able to evolve into what is commonly known as Gen AI; marking a significant shift in the AI landscape. One key evolution is GenAI's ability to create new data instances that are similar to the training data it ingested. This is achieved through techniques like generative adversarial networks (GANs) and deep learning, enabling AI to generate new content from user inputs such as images, music, or text.

**Diagram 1:** Predictive vs Generative AI

	Predictive AI	Generative AI (GenAI)
Functionality	Designed to analyze data to provide insights, make predictions, or make decisions based on the input it receives. Examples include recommendation systems and predictive analytics.	Designed to create new content. It can generate text, images, music, and more. Examples include AI that can write articles or create artwork.
Input Characteristics	Typically requires structured data to analyze. The success of the analysis largely depends on the quality and relevance of the input data.	Can work with a wider variety of data including both structured and unstructured data. It uses this data as a basis for generating new content.
Precision	Depends on the quality of the data it analyzes. With high-quality data, it can make highly accurate predictions or decisions.	Precision is not always the goal with GenAI. While it can be trained to generate high-quality outputs, the focus is often on creativity and novelty rather than precision.
Nature of Results	Typically, direct responses to the input data, such as a predicted outcome, a recommended action, or an insight into the data.	New creations that did not exist before. Results are derived from the input data but are not direct responses to it.

# Risk and Compliance in GenAI

Whether leveraging Predictive AI or GenAI, compliance over the technology remains at the forefront of the conversation. For an Information Technology (IT) compliance professional with experience in mainframe, on-prem, or cloud-based technologies, domains such as infrastructure security, data security, and data privacy remain high-focus areas. Further, evidence for compliance of these domains is universal for traditional technologies. For example, the compliance requirement “AWS S3 bucket is encrypted” is the same requirement applied across all S3 buckets in the control environment. While IT compliance is ever evolving, this form of “business as usual” compliance is predictable and manageable by IT compliance professionals. When considering GenAI compliance, one must consider traditional IT compliance challenges and focus areas (e.g., infrastructure security, data security, and data privacy) as well as challenges specific to GenAI, and the many models one organization may maintain. With GenAI, data issues are amplified, concerns over model security, model drift issues, and biased models supplement existing IT compliance considerations. Further, evidence collection becomes shades of grey. Rather than universal evidence across all S3 buckets, evidence requires context.

Traditional IT environments allow for standardized evidence collection and predictability; however, for AI models, the evidence can look different across each organization, and perhaps even within multiple models in the same organization. This additional layer of complexity multiplies the IT compliance professional’s surveillance of their control environment.

One of the largest concerns companies face in piloting their AI usage is a shift in its mindset. In a typical IT environment, an application undergoes testing in a non-production environment prior to ingesting true production or customer data. However, if the same mindset is applied to AI models, limitations to model functionality can be unveiled. GenAI models require a broad range of data for training. If a model is in a “non-production” environment, it is often difficult for companies to identify a data set that is suitable for testing in a non-production environment. It is not recommended to utilize sensitive or “production” data within a non-production environment, so test data is often “dummy data”. However, the quality of the model’s output is correlated with the quality of data it ingests. If unrealistic data is training the model, it

may pose issues for the reliability of the model’s output in the future. As such, compliance professionals should remain diligent in understanding the quality of data and the model’s environment that is in use for training. Cutting corners in either direction could cause potential hardships in the future in the form of security concerns or inaccurate models.





### **Amplification of Data Issues**

For many industries, data privacy continues to be a chief concern as companies try to solve the ever-present problem of 'human intervention,' which could translate to hacks and breaches internally or externally with trusted third-party service providers or erroneous data handling. These existing concerns around handling of sensitive data are amplified when using AI.

### **Model Security Concerns**

For the model to perform as designed, it is required to collect, process, store, and essentially learn from the data it ingests. The continuous and detailed data collection methods allow for the model to have access to a wide array of sensitive data, potentially leading to overexposure if the security measures are circumvented.

### **Model Drift Issues**

Model drift is a concept where, over time, the model's accuracy decreases due to a change in data patterns, behaviors, or collection methods that may require model re-training or transformation. Further, unless there is careful monitoring over the output, there may be a delay in realizing inaccurate outputs. Companies that rely on GenAI models for customer-facing or internal processes may experience inaccurate outputs from the model; resulting in potential customer, reputational, or compliance concerns.

### **Evidence of Bias**

A common misconception of AI models is that it removes human error and bias from its processing and outputs; however, one may argue that this phenomenon is even more present. The model is trained using existing human data and as a result, biases may be present. This is a heightened area of risk when considering the potential reputational and customer impacts companies face.

### **Applicable Controls to Address AI Risks**

Many resilient companies have found themselves at crossroads with evolutionary technology more than once and have emerged successful by centering the problem using a risk-based approach. By defining the risks, companies are able to create appropriate preventive, detective, and monitoring controls to mitigate the risk and reduce their exposure. To provide leading practices for businesses navigating AI using a risk-based approach, Deloitte has developed the Trustworthy AI risk framework which consists of 7 domains and 26 sub-components of risk and potential mitigating controls. The Trustworthy AI framework helps clients responsibly harness the power of AI for the benefit of shareholders and society at large. The Trustworthy AI framework provides a broad approach to helping clients implement sustainable, safe, and responsible AI solutions.

**Diagram 2:** Trustworthy AI™ is a framework to help effectively manage unique risks of AI

**Foundations of a responsible AI Risk Management Framework**

Comprehensive AI risk management principles serve as the cornerstone of a sound AI risk management framework. Deloitte’s Trustworthy AI™ framework provides the backdrop to a sustainable, safe AI risk management program.

**Safe/Secure**

AI systems can be protected from risks (including Cyber) that may cause physical and/or digital harm.

**Robust/Reliable**

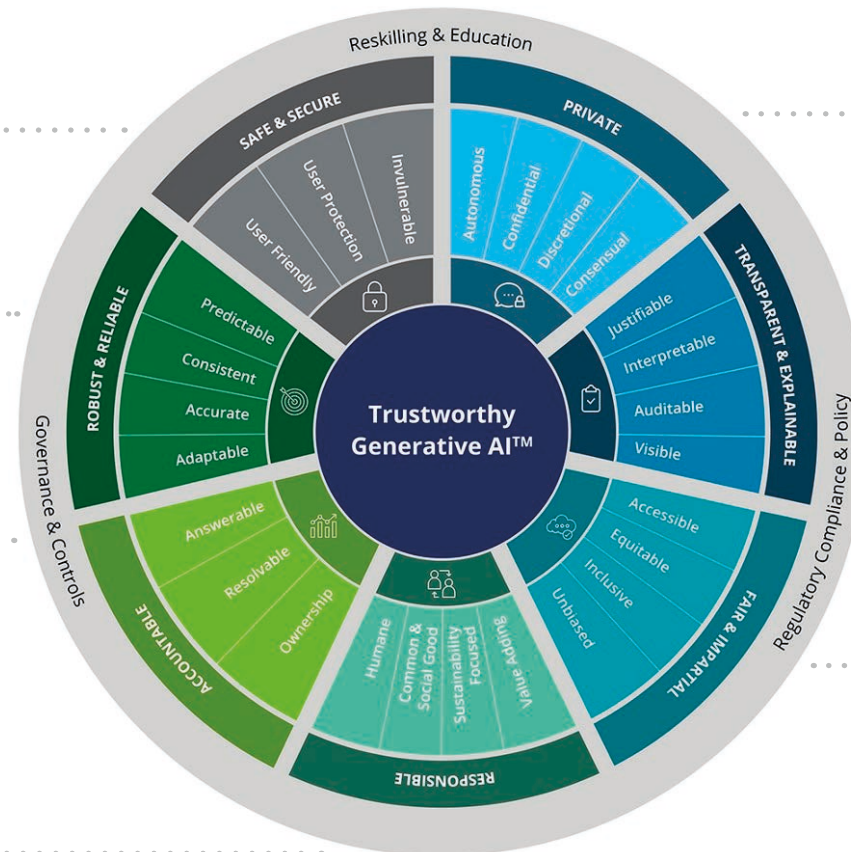
AI systems have the ability to learn from humans and other systems and produce consistent and reliable outputs.

**Accountable**

Policies are in place to determine who is responsible for the decisions made or derived with the use of technology.

**Responsible**

The technology is created and operated in a socially responsible manner.



**Private**

Consumer privacy is respected, and consumer data is not used beyond its intended and stated use; consumers are able to opt in/out of sharing their data.

**Transparent/Explainable**

Participants are able to understand how their data is being used and how AI systems make decisions; algorithms, attributes, and correlations are open to inspection.

**Fair/Impartial**

AI applications include internal and external checks to help ensure equitable application across participants.

### **Data Handling**

In our experience, one of the largest concern for organizations is the handling of data by AI models. Inappropriate handling of data and system logs may expose user data, personally identifiable information (PII), or system interactions that could jeopardize the confidentiality of data. To mitigate inappropriate handling of data, organizations could look to create an AI system cybersecurity architecture that includes considerations to demonstrate user information security and privacy is maintained. Periodic internal and external testing of the strength of the architecture to confirm the effectiveness of these controls should be performed. Organizations can look to identify potential opportunities for data minimization in relation to confidential data elements and explore strategies to reduce data collection, storage, and usage while making sure that the system's functionality and objectives are still met. Overall, a unified approach to protection should be developed to maintain the effectiveness and alignment of implemented controls.





### Security Concerns

IT vulnerabilities as a critical risk for IT compliance professionals are not a new concept in technology, and they remain a critical risk in AI models. Vulnerabilities in the system could allow malicious entities to gain inappropriate access, user data to be compromised, configurations and code changed for malicious purposes, or a loss in system availability. To help mitigate these risks, several controls can be implemented:

**Diagram 3:** Controls addressing vulnerabilities in the system



#### Continuous Monitoring

Continuous monitoring through system scanning for known and discovered vulnerabilities are remediated in an assigned timeframe.



#### Contingency Plan

Established contingency plan that outlines essential business functions associated with contingency requirements, including recovery and restoration timeframes.



#### Configuration Settings

Established configuration settings for the system are configured and documented.



#### Data Backup & Restore

Data backup and restore processes for the AI system including regular backup procedures, secure storage, and review of restore capabilities.



#### Access Management

Established account management process to easily create, modify, and remove users based on role-related criteria.



### **Model Drift**

A hallmark of the Trustworthy AI framework is to create AI models that are robust and reliable. An important factor of reliability is the measure of consistency in model output. The risk that the AI system performance drifts over time due to inconsistent data can cause output to be unreliable. Failure to monitor and diagnose AI model drift can lead to inaccurate predictions and potentially harmful consequences. To help mitigate this risk, controls to evaluate the AI system and model quality (sources, inputs, consistency/model drift) should be established and documented, including performing a risk assessment, establishing a cadence for review, and determining how identified issues will be addressed.

### **Biased Outputs**

In the context of fair and impartial AI, use of biased data in AI system training becomes ingrained in the system's algorithms and results in unfair, distorted, or discriminatory outputs. As the organization is using data from humans who may have biases, organizations could look to develop and implement controls to validate that data inputs are diverse, representative, and free from historical biases through defining non-discriminatory algorithm criteria and documenting comprehensive bias and exploratory data analysis. Achieving bias-free AI outputs requires recognition of the risk and appropriate mitigating actions.

### **Governance Approach for GenAI**

Organizations and IT compliance professionals are being asked to adapt once again to groundbreaking technology and modify or create an appropriate governance approach. As discussed, a risk-centered approach is critical to understanding key elements of the technology and adapting the organization's control environment to mitigate the identified risks. However, organizations and individuals should perform enhanced due diligence and look to outside sources to leverage such as regulatory guidance, industry-groups, Deloitte's Trustworthy AI framework, and vendor-provided guidance for AI tools.

# GenAI Compliance Landscape

Although there are well-established standards and frameworks specifically designed for AI, there is a noticeable absence of those tailored for GenAI. Predominant frameworks such as the National Institute of Standards and Technology (NIST) AI 100-1 (also known as AI Risk Management Framework 1.0)<sup>1</sup> and International Organization for Standardization (ISO)/ International Electrotechnical Commission (IEC) 42001<sup>2</sup> can provide companies with a foundation for initiating AI compliance. However, these require substantial time and ongoing updates to address the risks associated with GenAI applications. Regulatory bodies have started taking strides toward the regulation of GenAI. Notable examples include the United States (US) AI Executive Order (formally known as Executive Order 13859<sup>3</sup>, “Maintaining American

Leadership in Artificial Intelligence”), European Union Artificial Intelligence Act<sup>4</sup>, the European Commission’s Ethics Guidelines for Trustworthy AI<sup>5</sup>, the Montreal Declaration for Responsible Development of Artificial Intelligence<sup>6</sup>, the Organisation for Economic Co-operation and Development (OECD) Principles on AI<sup>7</sup>, and Singapore’s Model AI Governance Framework<sup>8</sup>, among others.

These initiatives signify an emerging awareness of the necessity for GenAI regulation. While these regulations represent substantial advancement toward a framework around GenAI compliance, the path forward remains complex given the lack of specific controls for companies to implement effectively.



<sup>1</sup> [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\) \(nist.gov\)](https://nist.gov/artificial-intelligence-risk-management-framework-ai-rmf-1.0)

<sup>2</sup> [ISO/IEC 42001:2023 - Artificial intelligence — Management system](https://www.iso.org/standard/72431.html)

<sup>3</sup> [Federal Register :: Maintaining American Leadership in Artificial Intelligence](https://www.federalregister.gov/documents/2023/03/13/2023-05130-artificial-intelligence)

<sup>4</sup> [Artificial intelligence act | Think Tank | European Parliament \(europa.eu\)](https://www.europa.eu/artificial-intelligence-act)

<sup>5</sup> [Ethics guidelines for trustworthy AI | Shaping Europe's digital future \(europa.eu\)](https://www.europa.eu/ethics-guidelines-for-trustworthy-ai)

<sup>6</sup> [declarationmontreal-iaresponsable.com/wp-content/uploads/2023/04/UdeM\\_Decl\\_IA-Resp\\_LA-Declaration-ENG\\_WEB\\_09-07-19.pdf](https://www.declarationmontreal-iaresponsable.com/wp-content/uploads/2023/04/UdeM_Decl_IA-Resp_LA-Declaration-ENG_WEB_09-07-19.pdf)

<sup>7</sup> [OECD Legal Instruments](https://www.oecd.org/legal/instruments/)

<sup>8</sup> [sgmodelaigovframework2.pdf \(pdpc.gov.sg\)](https://www.sgmodelaigovframework2.pdf)

# How Deloitte and AWS Can Help

To take the lead in compliance management, adopting a proactive strategy is essential. If you are an organization aiming to craft your own GenAI solutions, Amazon Bedrock offers built-in safeguards for a trusted deployment. The use of AWS Audit Manager can help your organization enhance its compliance even further with its GenAI-specific framework, produced through a collaboration between AWS and Deloitte.

To create a broad view of your organization's compliance posture beyond the cloud, Deloitte Nexus can effectively integrate AWS Audit Manager findings with other resources within your organization. This integrated approach not only helps your organization align with regulations and identify potential gaps in controls, but also empowers your organization to navigate the complexities of GenAI with confidence.





# Building within Guardrails using Amazon Bedrock






For companies looking to build their GenAI solution on the cloud, a suite of strategies and tools are available to assist organizations in remaining compliant while capitalizing on the latest technologies. Companies can build their GenAI applications on Amazon Bedrock, which offers a broad solution for building and scaling GenAI applications. Amazon Bedrock includes guardrails that are designed to help organizations deploy GenAI applications with trust. These guardrails provide a framework for maintaining safety, privacy, and compliance with responsible AI policies across various applications and use cases.

In addition to these guardrails, Amazon Bedrock supports the fine-tuning of models to increase accuracy for specific tasks, so customer data remains secure and private. Agents for Amazon Bedrock also allow for the execution of multistep tasks using company systems and data sources, simplifying the development and deployment of GenAI applications. These features collectively provide a robust framework for organizations to deploy GenAI applications with confidence, to help maintain compliance with responsible AI principles and enhancing user safety and privacy.





Diagram 4: Guardrails within Amazon Bedrock<sup>9</sup>

<b>Customized Safeguards</b> 	<b>Consistent AI Safety</b> 	<b>Content and Topic Filters</b> 	<b>PII Redaction</b> 	<b>Monitoring and Analysis</b> 
<ul style="list-style-type: none"> <li>Amazon Bedrock’s guardrails allow for the implementation of safeguards tailored to specific application requirements and responsible AI policies.</li> <li>This feature helps in promoting safe interactions between users and GenAI applications, aligning with company policies and principles to provide relevant and safe user experiences</li> </ul>	<ul style="list-style-type: none"> <li>The guardrails are designed to help you achieve a consistent level of AI safety across all your applications, regardless of the underlying foundation model (FM).</li> <li>They enable the evaluation of user inputs and FM responses based on use-case-specific policies, providing an additional layer of safeguards.</li> <li>You can create multiple guardrails with different combinations of controls and apply these across different applications and use cases</li> </ul>	<ul style="list-style-type: none"> <li>Organizations can define topics to avoid within the context of their application using natural language descriptions.</li> <li>Guardrails help identify and prevent user inputs and FM responses associated with restricted topics, providing relevant and secure user experiences.</li> <li>Additionally, content filters with adjustable thresholds provide an extra layer of control by filtering harmful content, including hate speech, insults, and violence, supplementing the protections integrated into the FMs.</li> </ul>	<ul style="list-style-type: none"> <li>Bedrock guardrails allow the detection of PII in user inputs and FM responses.</li> <li>This capability can enable organizations to either reject inputs containing PII or redact PII from FM responses, further protecting user privacy</li> </ul>	<ul style="list-style-type: none"> <li>Guardrails for Amazon Bedrock integrates with Amazon CloudWatch, allowing for the monitoring and analysis of user inputs and FM responses that violate policies defined in the guardrails.</li> </ul>

<sup>9</sup> [AI Safety - Guardrails for Amazon Bedrock - AWS](#)

# Adopting Gen AI Framework on AWS Audit Manager

AWS Audit Manager can be integrated with Amazon Bedrock to facilitate compliance processes. It provides a pre-established framework that helps organizations understand the performance of their GenAI implementation in relation to AWS's best practices. This framework allows Amazon Bedrock users to automate audits and evidence gathering, monitor AI model usage, identify sensitive data, and issue alerts on potential concerns. The 'Generative AI Best Practices Framework' is available in all AWS regions where Amazon Bedrock is offered, containing 110 new controls focused on governance, data security, incident management, and business continuity.<sup>10</sup> Organizations can start automated assessments by selecting desired controls and can

configure frameworks and data sources for tailored monitoring and reporting.

This framework was developed by AWS experts specializing in AI, compliance, and security assurance, together with Deloitte, which was AWS' Global GSI Security Partner of the Year for 2023. The framework is built upon industry frameworks such as NIST AI 100-1 and ISO/IEC 42001, Deloitte's Trustworthy AI framework, and various academic publications. These controls cover off on the following domains:

This framework serves as a compliance guide for organizations aiming to design and deploy GenAI

applications. It provides a broad view of AI system development, focused on key areas such as data preparation, model deployment, privacy, and risk management. The framework assimilates industry leading practices and emphasizes regulatory compliance and privacy protection, which can help enhance an organization's reputation and reduce data breach risks. It also focuses on risk management strategies for GenAI, making it a versatile resource for those in the GenAI landscape. The framework offers control measures and guidelines around accuracy, governance, and security in GenAI systems. It also stresses transparency, stakeholder communication, and ongoing monitoring to maintain system trustworthiness.

**Diagram 5:** Generative AI Best Practices Framework Domains



<sup>10</sup> [AWS Audit Manager introduces framework for generative AI on Amazon Bedrock](#)

# Extending Cloud Compliance Monitoring with Deloitte Nexus and AWS Audit Manager

Recognizing that organizations will have resources that are not solely within AWS, Deloitte has developed Deloitte Nexus, a suite of digital services designed to help businesses transform how they manage and optimize their risk and control programs to meet their regulatory and compliance needs. It combines advanced technology, like Deloitte Nexus Digital Nerve Center™ platform, with a highly skilled workforce specializing in data science, engineering, and cloud regulatory risk

and control specialists that can tailor your needs into a single pane of glass to help the organization understand whether it is meeting regulatory, legal, and compliance requirements. As a trusted advisor, helping organizations adopt cloud, our governance, risk, and compliance and data analytics specialists can simultaneously gather the technical requirements as you adopt cloud to automatically gather the data requirements to align to your cloud risk and controls framework. The Deloitte Nexus framework

leverages three foundational pillars to operationalize your transformation: Command Center, Digital Risk & Controls, and Workforce Transformation. It provides near real-time intelligence and transparency into the operational health and risk posture of a business. It can help enhance risk management with elevated risk acumen and insights utilizing machine learning and artificial intelligence capabilities.

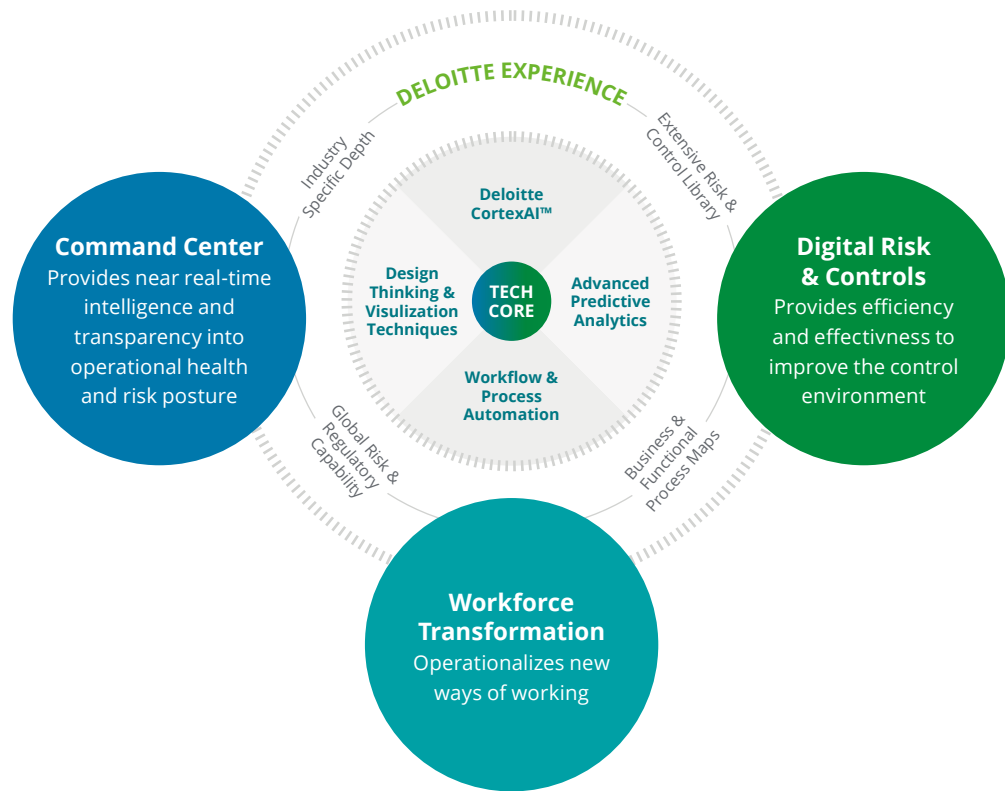


Diagram 6: Deloitte Nexus Digital Nerve Center™ platform

**Keep your finger on the pulse of your business**

Nexus is a suite of innovative digital services designed to help you transform the management and optimization of your organization’s risk and control programs.

Nexus is designed to help you transform the way your organization manages and optimizes your business processes, risks, and controls through our digital nerve center and can provide the following potential outcomes:



**Enhanced Risk Management**

Elevate risk acumen and insights utilizing machine learning (ML) and artificial intelligence (AI)



**Risk Sensing**

Provide insight and foresight prior to issues happening to help improve resiliency



**Workforce Elevation**

Increase professional productivity while focusing on exceptions and core risk management



**Organizational Agility**

Build elasticity in your workforce through management of digital and co-sourced FTEs



**Cost savings**

Drive better efficiency through decreasing manual processes and providing rapid results



**Governance & Oversight**

Promote greater sense of ownership, visibility and accountability across the organization



**Testing and Monitoring**

Modernize outdated controls evaluation programs while improving quality



**Customer Experience**

Enhance delivery by helping to mitigate operational issues and improve processing, creating an enhanced experience for your customer












Diagram 7: Deloitte Nexus Digital Nerve Center™ integration with Audit Manager

### Streamlining Audit & Compliance with Nexus Digital Nerve Center Integration

With Deloitte’s extensive experience with risk and compliance, we can design and deliver an end-to-end and continuous auditing and compliance solution using Nexus alongside AWS Audit Manager to improve audit efficiency and easily scale over rapid cloud adoption.

Capabilities	Assessment objective & scope 	Assessment initiation 	Evidence collection 	Control analysis 	Assessment reporting 
 <b>Audit Manager Capabilities</b>	<ul style="list-style-type: none"> <li>Provide a list of AWS resources and services that can be selected for the assessment</li> </ul>	<ul style="list-style-type: none"> <li>Provide pre-built assessment frameworks</li> <li>Ability to customize frameworks or add custom controls for the assessment</li> <li>Ability to delegate or assign control ownership to individuals</li> </ul>	<ul style="list-style-type: none"> <li>Supports automated collection of evidence daily</li> <li>Preserve integrity of gathered evidence</li> <li>Capture evidence metadata including source, timestamp, and status</li> </ul>	<ul style="list-style-type: none"> <li>Convert collected evidence into a consistent auditor friendly report</li> <li>Assist in evidence assertion on some controls (Manual review of evidence for control assertion may be required)</li> </ul>	<ul style="list-style-type: none"> <li>Reports can be generated and accessed by relevant teams/personnel</li> </ul>
 <b>How can Nexus help?</b>	<ul style="list-style-type: none"> <li>Develop a customized cloud control framework</li> <li>Scope AWS accounts, services, and resources using a risk-based approach</li> <li><b>Include external technology elements (outside AWS) within the scope of testing</b></li> </ul>	<ul style="list-style-type: none"> <li>Configure Audit Manager based on the audit program</li> <li>Conduct test runs to validate cost expectations</li> <li>Establish a continuous auditing and compliance program</li> </ul>	<ul style="list-style-type: none"> <li>Review a sample of evidence to validate completeness and accuracy</li> <li>Operate a continuous auditing and compliance program</li> <li><b>Aggregate data from different external sources (outside AWS) for analysis</b></li> </ul>	<ul style="list-style-type: none"> <li>Conduct control testing for controls not evaluated by Audit Manager</li> <li><b>Generate AI Logic to determine compliance with selected scope</b></li> </ul>	<ul style="list-style-type: none"> <li>Real-time and continuous feedback on compliance excellence</li> <li>Visualize metrics and trends (AI capability)</li> <li><b>Provide custom reporting Dashboards that suit the needs of stakeholders</b></li> <li><b>Perform trend analysis to determine improvement areas</b></li> </ul>

# Efficiently Navigating the GenAI Compliance Landscape: A Broad Approach

The comparison between Predictive AI and Generative AI underscores the specific considerations and challenges inherent to each, particularly in the realm of compliance. Whether it's the amplification of data issues, model security concerns, model drift issues, or the need for contextual measures to prevent bias, both forms of AI present distinct compliance considerations.

However, these challenges can be navigated with the right tools and strategies. Deloitte's Trustworthy AI framework, Amazon Bedrock's built-in safeguards, and the integration of AWS Audit Manager findings via Deloitte Nexus—when combined—provide a broad approach to managing GenAI compliance.

The application of these tools and strategies, combined with a firm understanding of the unique nature of GenAI, can help in effectively addressing data handling concerns, biased outputs, and security issues.

Ultimately, as the GenAI landscape continues to evolve, so too must the strategies for managing its associated risks. By adopting a proactive approach and leveraging the right tools and frameworks, organizations can harness the transformative power of GenAI while maintaining a strong compliance posture.



# Authors




**Charlie Willis**  
**Managing Director**  
Digital Controls  
Cloud Digital Controls Leader  
Deloitte & Touche LLP  
chwillis@deloitte.com




**Christina De Jong**  
**Partner**  
Technology Media Telecom  
Regulatory Leader  
AWS Lead Business Partner  
Deloitte & Touche LLP  
christinadejong@deloitte.com



**Casey Kacirek**  
**Managing Director**  
IT & Specialized Assurance  
Trustworthy AI™ Leader  
Deloitte & Touche LLP  
ckacirek@deloitte.com



**Alex Vorpahl**  
**Senior Manager**  
Internal Audit  
Deloitte Nexus  
Deloitte & Touche LLP  
avorpahl@deloitte.com



**Elaine Li**  
**Manager**  
AWS Cyber  
Cloud Digital Controls  
Deloitte & Touche LLP  
elaineli2@deloitte.com



**Lauren Goodchild**  
**Senior Consultant**  
Digital Controls  
Deloitte & Touche LLP  
lgoodchild@deloitte.com

---

## Special Thanks To

**Neha Singh Rajpurohit**  
**Senior Product Manager – Technical**  
AWS Audit Manager  
nrsingh@amazon.com

**John Fischer**  
**Senior Consultant**  
AWS Audit Manager  
jofisc@amazon.com





This document contains general information only and Deloitte is not, by means of this document, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This document 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 advisor.

Deloitte shall not be responsible for any loss sustained by any person who relies on this document.

All product names mentioned in this document are the trademarks or registered trademarks of their respective owners and are mentioned for identification purposes only. Deloitte & Touche LLP is not responsible for the functionality or technology related to the vendor or other systems or technologies as defined in this document.

As used in this document, "Deloitte" means Deloitte & Touche LLP, a subsidiary of Deloitte LLP. Please see [www.deloitte.com/us/about](http://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.

Designed by CoRe Creative Services. RITM1764403