

THE IMPACT OF AI ON INTERNAL AUDIT

BCS IRMA Conference

June 2025

protiviti®
Global Business Consulting

Agenda

AI Trends

Internal Audit Use Cases

Risks and Regulations

Questions

Introductions

Alex Psarras



Director

NextGen Internal
Audit UK Leader

Protiviti



AI Trends

Some Definitions

Specific definitions and distinctions do not have academic and scientific consensus. For practical application, the broad term “AI-ML” is used to describe systems that emulate human thought without explicit programming.

Machine Learning (“ML”)

The ability of algorithms and statistical models to learn and adapt without being explicitly programmed

Large Language Models (“LLMs”)

A type of generative AI trained on vast amounts of data with a large number of parameters that generate novel text-based responses



Artificial Intelligence (“AI”)

The ability of a computer system to emulate human-like cognitive abilities, such as learning and problem-solving

Generative AI (“GenAI”)

A type of AI capable of creating new and novel outputs, such as text, code, images, videos, and audio

Agentic AI

Systems that perceive their environment, set or pursue objectives and autonomously take actions via planning and evaluation loops.

Adoption trends being observed

- 01 Generative AI Maturity**
Transition from Hype to Reality: After the initial excitement surrounding generative AI, organizations are moving from experimentation to practical implementation. This shift is characterized by a more nuanced understanding of AI capabilities and limitations, as businesses seek to integrate AI solutions at scale
- 02 Integration into Existing Systems**
Complementary Tools: Many organizations are embedding AI into existing business applications rather than replacing them. Features like Microsoft's Copilot and Adobe's Generative Fill illustrate how AI can enhance traditional workflows without overhauling established systems
- 03 Cost Reduction and Automation**
Driving Factors: The primary motivations for AI adoption include the need to reduce operational costs and automate processes. Companies are increasingly leveraging AI tools that make these capabilities more accessible, thereby improving efficiency across various functions
- 04 Focus on Multimodal Models**
Diverse Data Inputs: The next wave of AI advancements is expected to focus on multimodal models that can process and analyze various types of data simultaneously. This development aims to enhance the versatility and applicability of AI solutions across different sectors
- 05 Ethics and Governance**
Regulatory Considerations: Organizations are increasingly aware of the ethical implications of AI, leading to a focus on governance and compliance. The evolving regulatory landscape is prompting businesses to adopt more responsible AI practices and ensure that their models are free from bias

Key Trends

AI market is expected to reach


\$17.7 trillion by **2030**

ChatGPT had **1 million** users within the first five days of being available

64% of businesses expect AI to increase productivity

AI is expected to see an annual growth rate of **37.3%** from 2023 to 2030

The Large Language Model Performance Race



Humanity's Last Exam Benchmark

Model	Accuracy (%)
o3	20.3
Gemini 2.5 Pro	18.4
o4-mini	18.1
DeepSeek-R1-0528	17.7
o3-mini	13.4
Gemini 2.5 Flash	12.1
Qwen3-235B	11.8
Claude 4 Opus	10.7
DeepSeek-R1	8.5
Claude 3.7 Sonnet	8
o1	8

Classics


Question:



Here is a representation of a Roman inscription, originally found on a tombstone. Provide a translation for the Palmyrene script. A transliteration of the text is provided: RGYN^o BT HRY BR ^cT^o HBL

Chemistry

Question:



The reaction shown is a thermal pericyclic cascade that converts the starting heptaene into endiandric acid B methyl ester. The cascade involves three steps: two electrocyclizations followed by a cycloaddition. What types of electrocyclizations are involved in step 1 and step 2, and what type of cycloaddition is involved in step 3?

Humans and AI: Do they work better together or alone?

by MIT Sloan Office of Communications | Oct 28, 2024

Neuralink brain chip helps first paralysed patient play chess using 'telepathy'

'It has already changed my life,' says first human trial patient of Elon Musk's brain-computer interface

Original Investigation | Health Informatics

July 16, 2024

Large Language Model-Based Responses to Patients' In-Basket Messages

Nvidia CEO predicts the death of coding — Jensen Huang says AI will do the work, so kids don't need to learn

Anthropic boss: AI will take half of entry level jobs in the UK


BUSINESS

The CEO of Zoom wants AI clones in meetings

Zoom founder Eric Yuan has big ambitions in enterprise software, including letting your AI-powered 'digital twins' attend meetings for you.

protiviti®

Amazon's Just Walk Out stores relied on '1,000 people in India watching,' not AI



Reuters > Artificial Intelligence

OpenAI and others seek new path to smarter AI as current methods hit limitations

OpenAI software ignores explicit instruction to switch off

ChatGPT maker's 'most capable' model sabotages shutdown mechanism



AI model threatened to blackmail engineer over affair when told it was being replaced: safety report

Microsoft Recall Gets Recalled

Microsoft's Copilot Plus PCs face backlash over the Recall feature, raising privacy concerns and prompting a quick recall for revisions.

Warning social media videos could be exploited by scammers to clone voices

Research released by Starling Bank finds 28% of people have been targeted at least once in the past year

The Fake Drake AI Song Earned Millions of Streams – But Will Anyone Get Paid?


[ft.com](https://www.ft.com/content/00000000-0000-0000-0000-000000000000) / OpenAI

Scarlett Johansson hits out at ‘eerily similar’ OpenAI chatbot voice

What is Artificial General Intelligence (AGI)?

	Artificial Narrow Intelligence	Artificial General Intelligence
Definition	Designed to perform a specific task or a limited set of tasks.	Aims to replicate human-like intelligence, capable of performing any intellectual task a human can do.
Focus	Specific task-focused	Broad and generic
Scope	Limited to one domain	Cross-domain adaptability
Autonomy	Pre-defined functions	Independent, human-level reasoning
Abilities	Mimics one aspect of cognition	Replicates human-level cognition
Risk	Low (task-specific risks)	Moderate (alignment with human goals)

The Road to AGI



Reasoning AI

- Chain of Thought & Tree of Thought Reasoning
- Reinforcement Learning
- Incremental Problem Solving
- Self-Correction
- Safety and Alignment

Chain-of-Thought Prompting

Model Input


Q: Roger has 5 tennis balls. He buys 2 more cans of tennis balls. Each can has 3 tennis balls. How many tennis balls does he have now?

A: Roger started with 5 balls. 2 cans of 3 tennis balls each is 6 tennis balls. $5 + 6 = 11$. The answer is 11.

Q: The cafeteria had 23 apples. If they used 20 to make lunch and bought 6 more, how many apples do they have?

Model Output

A: The cafeteria had 23 apples originally. They used 20 to make lunch. So they had $23 - 20 = 3$. They bought 6 more apples, so they have $3 + 6 = 9$. The answer is 9. ✓



(d) Tree of Thoughts (ToT)



oyfjdnisdr rtqwainr acxz mynzbhhx -> Think step by step

Use the example above to decode: oyekaijzdf aaptcg suaokybhai ouow aqht mynznvaatzacdfoulxxz



Thought for 5 seconds

We are given:

First, an example:

"oyfjdnisdr rtqwainr acxz mynzbhhx -> Think step by step"

We can see that when "oyfjdnisdr rtqwainr acxz mynzbhhx" is transformed to "Think step by step"


Our task is: Use the example above to decode: "oyekaijzdf aaptcg suaokybhai ouow aqht mynznvaatzacdfoulxxz"

So the first part is to figure out how "oyfjdnisdr rtqwainr acxz mynzbhhx" is decoded into "Think step by step"

First discovery - Pattern Analysis:

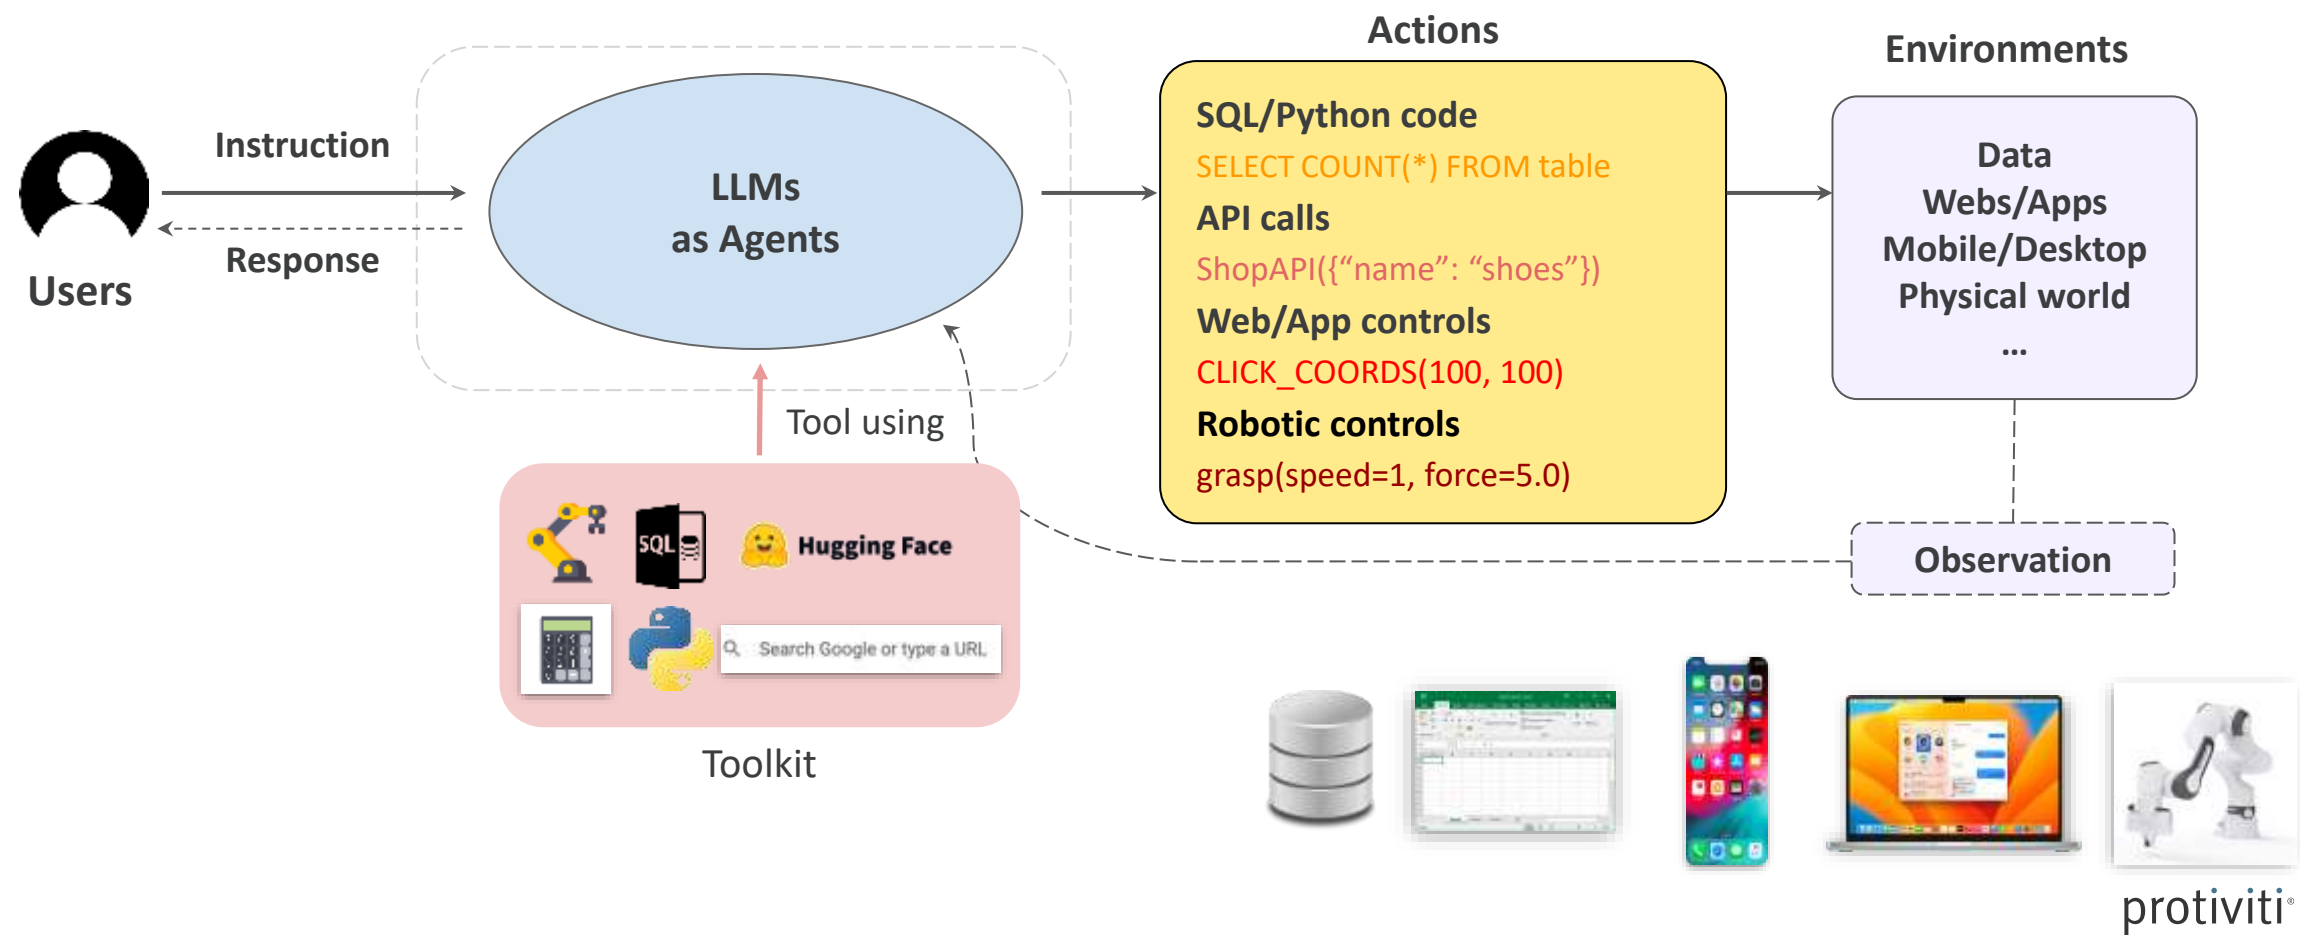
1. Count the letters in each word:


The Road to AGI



Large Language Models as Agents

Source: OSWorld: Benchmarking Multimodal Agents for Open-Ended Tasks in Real Computer Environments, The University of Hong Kong





Internal Audit Use Cases

Artificial Intelligence In Internal Audit



Artificial Intelligence

Enabling Technologies


Audit Functions leverage Machine Learning and Artificial Intelligence for the following benefits

Generate natural language from a machine-representation system such as a knowledge base or a logical form

Develop an action list by recording the user perform that task in the application's graphical user interface (GUI), and then perform automation by repeating those tasks

Gain high-level understanding from digital images or videos and seek to automate tasks that a human visual system could do

Use in activities such as inventory counts to determine whether or not image data contains some specific object, feature, or activity and identify the correct object, feature or activity, if such exists



Methodologies and technologies that enable the recognition and translation of spoken language into text by computers, such as technology used in virtual assistants

Process and analyse large amounts of natural language data, including natural language understanding, such as reviewing contract terms

Software agents that can learn or use knowledge to achieve their goals. They may be very simple or complex. Also, more than one agent can be trained in a competitive environment


Entails all aspects of designing, building and managing the automated decision-making systems that an organisation uses to manage its interactions with customers, employees and suppliers

Microsoft's Copilot Products

Copilot is a brand; a collection of tools and applications leveraging Generative Artificial Intelligence capabilities.


Microsoft has organised and embedded Copilot across multiple Microsoft products.

Most organisations start with Copilot for M365, which is the suite of Office applications.



Enabling Process Walkthroughs using Copilot in MS Teams

Leverage Copilot to not only capture and summarise meeting notes, but also steer discussions in real-time



AI-enabled Walkthrough Best Practices

Before the meeting:

- Prepare targeted questions related to the processes, controls, risks, and pain points to be discussed

During the meeting:

- Generate follow-up questions based on the discussion to dig deeper into potential issues or areas requiring clarification
- Reconcile the agenda against the transcript to ensure key points are not missed

After the meeting:

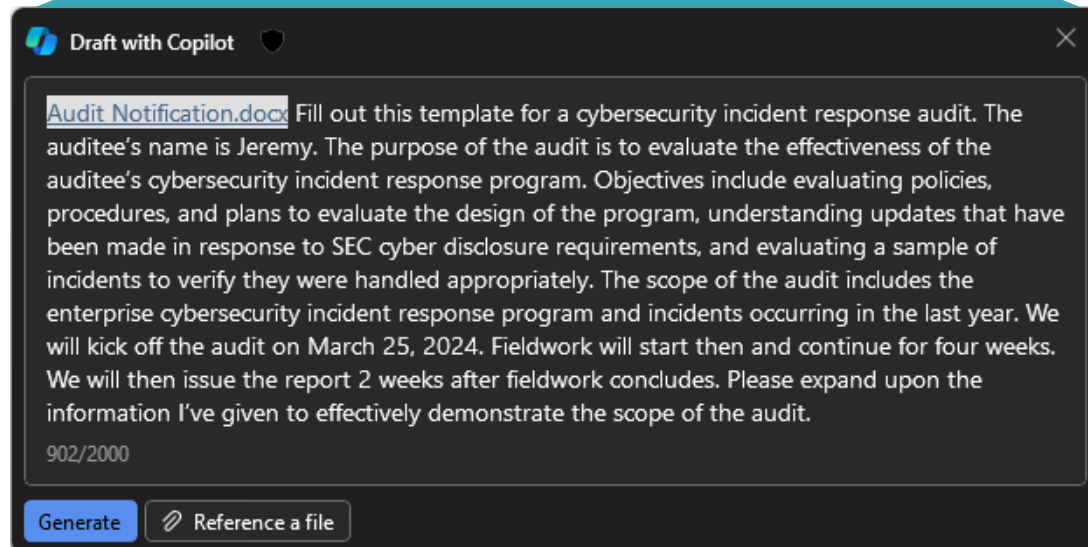

- Summarise the discussion, including key action items and follow-up requests
- Identify any potential control gaps or areas for improvement based on the AI's analysis of the meeting content



Important: Always obtain explicit permission from all participants before recording and be transparent about the intended use of AI to assist with note-taking, analysis, and summarisation to ensure trust and compliance with data privacy policies.

Populating the Audit Notification Template using Copilot for MS Word

It's as simple as opening Word, entering a prompt, and generating



Dear Jeremy,

As part of the internal audit plan, we are conducting a cybersecurity incident response audit with the purpose of evaluating the effectiveness of your cybersecurity incident response program. This audit is designed to assess the adequacy and compliance of your policies, procedures, and plans to respond to cyber incidents, understand the updates that you have made in response to the SEC cyber disclosure requirements, and verify that a sample of incidents that occurred in the last year were handled appropriately and in accordance with best practices.

In this email, we would like to share with you some preliminary information about the objectives, scope, and timeline of the audit. Please note that these are subject to change as the audit progresses and new information emerges.

Preliminary Objectives

- Evaluate the design and implementation of the cybersecurity incident response program, including roles and responsibilities, governance, reporting, escalation, and communication.
- Assess the alignment of the cybersecurity incident response program with the SEC cyber disclosure requirements and the NIST Cybersecurity Framework.
- Review a sample of cyber incidents that occurred in the last year and evaluate the effectiveness and efficiency of the incident response process, including identification, containment, eradication, recovery, and lessons learned.
- Identify any gaps, weaknesses, or areas for improvement in the cybersecurity incident response program and provide recommendations to enhance the program.

Preliminary Scope

- The enterprise cybersecurity incident response program, including policies, procedures, and plans.
- A sample of cyber incidents that occurred in the last year, covering different types, sources, and impacts of incidents.
- The documentation and evidence related to the incident response process, such as incident reports, logs, emails, action plans, and remediation steps.


Preliminary Timeline

- Kickoff: March 25, 2024
- Fieldwork: March 25, 2024 - April 22, 2024
- Report issued: May 6, 2024

The next step is to schedule a kickoff meeting with you and your team to discuss the audit process, expectations, and logistics. We will send you a meeting invite shortly.

We look forward to working with you on this audit and appreciate your cooperation and support. If you have any questions or concerns, please do not hesitate to contact us at <sender email> or <sender phone>.

Sincerely,



Analysing Data in Microsoft Excel

Open Excel, Format Data as a Table, Ask Questions

The screenshot displays the Microsoft Excel interface with the 'Table Design' tab active. A data table is visible with the following columns: User Name, Created, Deleted, Modified, Privileged?, and Group/User. The table contains 13 rows of data. To the right of the table, the 'Copilot (Preview)' pane is open, showing a query: 'how many enabled users have the domain admins group?'. The Copilot response indicates that there are 13 rows for 'Group Name' 'Domain Admins' and 'Enabled' 'True' and 'Group/User' 'User'. The Copilot pane also includes a button to 'Add to a new sheet' and a disclaimer: 'AI-generated content may be incorrect'.

User Name	Created	Deleted	Modified	Privileged?	Group/User
Administrator	#####		#####		User
Administrator	#####		#####	TRUE	User
Administrator	#####		#####		User
Administrator	#####		#####	TRUE	User
Administrator	#####		#####		User
Jerry Seinfeld	#####		#####		User
Guest	#####		#####		User
Kramer	#####		#####		User
SAA Elaine Benes	#####		#####		User
SAA Elaine Benes	#####		#####		User
SAA Elaine Benes	#####		#####		User
SAA Elaine Benes	#####		#####		User
SAA Elaine Benes	#####		#####		User
SAA Elaine Benes	#####		#####		User

Copilot (Preview)

how many enabled users have the domain admins group?

Number of rows for 'Group Name' 'Domain Admins' and 'Enabled' 'True' and 'Group/User' 'User'

Count of Rows

13

+ Add to a new sheet

AI-generated content may be incorrect



Today

ProGPT-Chat



GPTs

Explore Protiviti's custom GPTs built upon our Intellectual Property

Top CAE CIO CFO CRO

Planning	Fieldwork	Reporting	Other Internal Audit Use Cases	
 Audit Announcement	 Business Process Flows	 Executive Summary Development	 Internal Audit Charter Creation	 Audit Strategy & Roadmap
 Planning & Scoping Memo	 Process/Control Narratives	 Observations & Recommendations	 Audit Policies & Procedures	 Risk Assessment
 Document Request List	 Document Review Acceleration	 Audit Response Analyzer	 Audit Analytics & Automation	 Stakeholder Communication
 Work Program Development	 Work Program Completion	 Best Practices & Value Adds	 Meeting Notes Cleaner	 Function Research

Search GPTs

By using ProGPT, you agree to protect sensitive data, not store it in chat, and allow chat history collection for service improvements. This tool is for internal use only.

Meeting Transcript to Risk & Control Matrix and Process Map

ProGPT

Submit Feedback

New Chat

Document Chat

Documents

Persons

Today

SAP Access Proc...

Summarizes the meeting with key action items

4:25 PM

Initial Meeting Transcript.pdf

027 KB

Success: memory established

Summarize

4:33 PM

You are an internal auditor reviewing a meeting transcript of a process walkthrough. The attached transcript describes the process taken for provisioning new and modified access to SAP. Using this information, please create a Risk & Control Matrix (RCM) in table format that identifies and lists out the controls described in the walkthrough. Each control should be in a unique row and the 5 columns are: Control ID, Control Description, Key Risk, the key mitigate, Control Description: a more detailed description of the control, and the evidence. Please list the evidence: A list of evidence items needed to perform within each table cell.

ProGPT Chat

4:33 PM


Documents Referenced

Initial Meeting Transcript.pdf


AI generated content may be incorrect

Generating bot response...

The maximum size of the document is 10 MB, and only PDF and TXT files can be uploaded at this time.



Generating a First Draft Audit Report with the Click of a Button

The image is a screenshot of a web application interface for the 'Protiviti AI Audit Report Accelerator'. The header features a home icon, the text 'protiviti | Technology Accelerator Platform', and the product name 'Protiviti AI Audit Report Accelerator'. The main content area is titled 'Overview' and describes the tool's function: 'The Protiviti AI Audit Report Tool Accelerator, powered by ProtivitiGPT's LLM, swiftly crafts quality reports from your audit files. This tool is designed to elevate report quality and slash drafting time.' Below this is a section titled 'Streamlined Reporting' which states 'Expect concise, comprehensive reports with:' followed by a bulleted list: '• Executive Summary: Key insights at a glance.' and '• Detailed Findings & Recommendations: In-depth analysis and actionable steps.' The next section is 'Before Getting Started', which says 'Have these documents on hand with the following information:' followed by a numbered list: '1. Scoping Memo: background on the audit report, objectives, scope, and time period', '2. Work Program / RCM: conclusions on the effectiveness of each control', and '3. Observations Log: every observation noted throughout the audit project, along with risk rating'. A note follows: 'Note: The LLM cannot process graphical text, so please avoid submitting files where most of the information is in the form of pictures or diagrams.' The final paragraph states: 'Your data stays safe in our secure IT environment. Review Protiviti's usage guidelines and client contract AI-use provisions before starting'.

AI-Generated Video Audit Reports

VIDEO STARTS IN

10

Please Stand By

Agentic AI for Contract Reviews

Audit Background

Our organisation's central contract management repository contains **over 3,000 vendor agreements** in Word and PDF format. Following the rollout of enhanced Terms and Conditions and updated legal provisions, Internal Audit will be performing an end-to-end compliance review.

Audit Objective

To verify that all vendor agreements conform to the newly approved legal requirements and group-wide contractual standards, and to determine whether existing contracts require amendment or formal renegotiation.

Methodology


Rather than relying on a traditional statistical sampling framework, we will deploy autonomous AI agents to perform a full-population analysis of every contract document. This AI-driven approach will:

1. Extract critical clauses (e.g. indemnities, termination rights, data-protection obligations) from each agreement.
2. Compare extracted clauses against the updated legal requirements.
3. Flag discrepancies and classify them by severity and recommend next steps.
4. Compile findings into a single, comprehensive compliance dashboard and narrative report.

Risks of AI adoption

AI Risks

While the adoption of AI brings many advantages to many organisations, it also comes with several risks and potential challenges. Here are few of them:



Guardrails we see being adopted

Guardrails are being adopted in AI to ensure safe, effective, and trustworthy AI deployment. Guardrails help organizations navigate the complexities of GenAI, safeguarding against risks while maximizing the technology's transformative potential. Some examples of guardrails include

Data Management & Security

- Policies and proven techniques to protect proprietary and sensitive personal information.
- Advanced Data Management strategy

Knowledge Models, Graphs and RAG

- Ground the LLM for use in a specific domain.
- Increase the precision and accuracy of its output.



Prompt Engineering

- Use of Prompt Engineering.
- Reduce the possibility of behavior hijacking or poor-quality requests.

Flow Control and Analytics


- Quality control and regression checks.
- Embed the calls to LLMs in articulated and potentially complex workflows that check, enrich, filter and process the input before and possibly instead of calling the LLM.

Ethical AI guardrails ensure responsible innovation

Recognised EMERGING Standards for AI Governance

NIST AI Risk Management Framework is designed to equip organisations and individuals with approaches that increase the trustworthiness of AI systems, and to help foster the responsible design, development, deployment, and use of AI systems over time.

ISO/IEC AI Framework provides guidance on managing risk associated with the development and use of AI. The document offers strategic guidance to organisations to assist in integrating risk management into significant activities and functions.



Predominant Themes	
Safety, Security & Resilience	Transparency & Accountability
Privacy	Ethics & Fairness
Validity / Reliability	Explainability & Interpretability

Data Quality & Integrity

European Union AI Act focuses primarily on strengthening rules around data quality, transparency, human oversight and accountability. It also aims to address ethical questions and implementation challenges.

Generative AI-Specific Controls

Controls.....				
1 AI governance framework: Establish a comprehensive AI governance framework that includes policies, procedures, and guidelines to define the appropriate use of AI-generated content within the organization.	2 Human oversight and review: Implement a process that ensures human oversight and review of AI-generated content before it is disseminated, in order to identify and correct potential inaccuracies, biases, or inappropriate content.	3 Access controls and monitoring: Implement strong access controls to limit the use of AI-generated content to authorized personnel only, and monitor usage to detect and prevent potential misuse.	4 Data privacy and security measures: Establish robust data protection measures, including encryption, access controls, and secure storage, to safeguard the data used by AI systems and ensure compliance with relevant regulations.	5 Bias and ethical considerations: Implement a process for assessing potential biases and ethical issues related to AI-generated content and incorporate mechanisms to address them, such as bias detection and mitigation techniques.
6 Employee training and awareness: Develop and deliver training and awareness programs for employees to educate them on the appropriate use, limitations, and potential risks associated with AI-generated content.	7 AI system maintenance and updates: Regularly update and maintain AI systems to ensure optimal performance and address any identified issues, such as biases or inaccuracies, in the generated content.	8 Incident response and reporting: Establish an incident response process to handle any cases of AI-generated content misuse, including reporting mechanisms, investigation procedures, and corrective actions.	9 Performance measurement and continuous improvement: Measure the effectiveness of AI-generated content and related controls and use this information to drive continuous improvement of AI systems, policies, and procedures.	10 Collaboration with AI vendors: Maintain open communication and collaboration with AI vendors, Trusted Partner to stay informed of new features, best practices, and potential risks associated with the use of their technology.

Questions

Alex Psarras



Director

NextGen Internal
Audit UK Leader

Protiviti



Face the Future with Confidence[®]

© 2025 Protiviti – Confidential. Protiviti is not licensed or registered as a public accounting firm and does not issue opinions on financial statements or offer attestation services. All registered trademarks are the property of their respective owners.

protiviti[®]
Global Business Consulting