

THE POWER OF BEING UNDERSTOOD
ASSURANCE | TAX | CONSULTING

Ani

One of the
RSM team


RSM

AI governance unveiled:
Leveraging AI effectively and
responsibly

Presenters



Dave Mahoney

Director, Risk Consulting

National Artificial Intelligence-Risk Leader
CISSP



William Jones

Manager – Modern Work

William Jones is the operational lead for the Enterprise Content Management team at RSM US LLP. His team works with clients to implement Microsoft Copilot in their own environments, with a special focus on both Data Security and Copilot Adoption.

What is artificial intelligence?



Artificial Intelligence

Artificial intelligence is the development of systems to mimic human problem-solving behavior by computing a prediction and decision and then executing an action



Machine Learning

Machine learning is an artificial intelligence component that uses data and algorithms to imitate human learning.

The machine learns independently by drawing inferences from patterns in the data, gradually adapting and becoming more accurate.



Deep Learning

Deep learning is a type of machine learning that attempts to simulate the human brain by creating artificial neural networks that utilize and extract more information from the data.

GENERATIVE AI

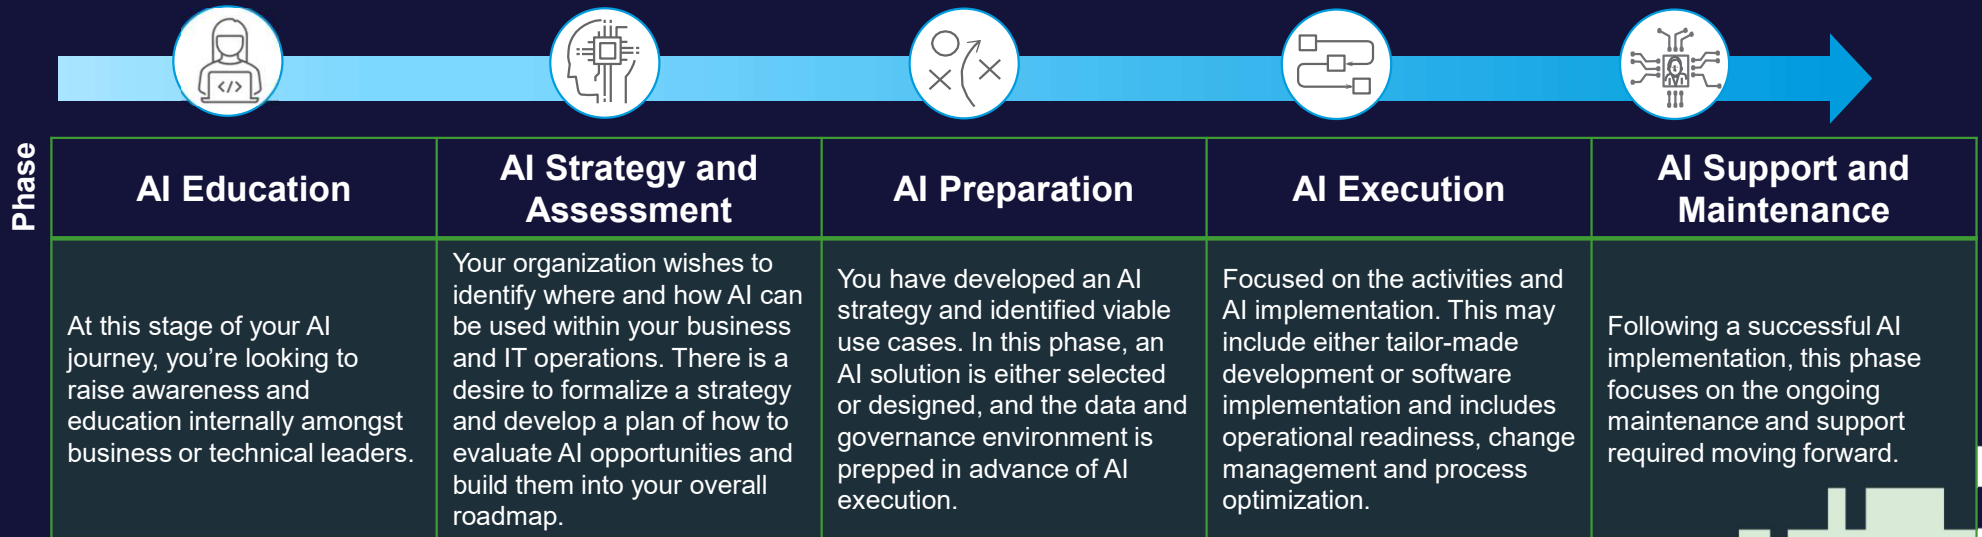
Generative AI is a type of deep learning algorithm that can generate content by learning & imitating the patterns and structure of data it was trained on.

LARGE LANGUAGE MODELS

Large language models are deep learning algorithms trained on vast datasets to understand, predict, and generate text in a human-like way.

GenAI Journey

AI customer journey outlines a strategic pathway for clients to harness the potential of AI at all phases. RSM can help you navigate this journey from initial education to implementing tailored AI solutions. We support our clients with successful adoption and the ability to make impactful decisions and ultimately become AI champions within your organization. Specific offerings by customer journey phase are outlined in the following slide.



What can AI do?

Artificial Intelligence (AI) harnesses advanced analytical and logic-based methods, such as machine learning (ML), to understand events, automate decisions, and initiate actions. We categorize AI's fundamental capabilities into three main domains to transform operations and drive innovation:



Trend Analysis

These applications identify patterns, classify information and make predictions based on existing data. This provides a birds eye view of business performance and provides the ability to anticipate future outcomes.

- **Pattern recognition identifies trends within large datasets and is adept at spotting irregularities that may indicate fraud, operational hiccups, or security breaches**
- **Enhanced classification by categorizing data into meaningful segments**
- **Predictive capabilities by leveraging historical data to forecast business metrics**



Workflow Automation

By integrating intelligent automation into workflows, companies can drastically reduce manual efforts and costs. This allows focus to redirect on innovation and strategic tasks that AI cannot replicate, increasing overall productivity.

- **Automates routine tasks, increasing efficiency.**
- **Processes documents quickly, cutting down on manual work.**
- **Supports complex tasks with data-driven insights.**
- **Optimizes resource use for better outcomes.**



Generative & Autonomous Technologies

These systems interact with their environment, users, or operate independently. They offer ways to enhance customer engagement, automate logistics, and innovate product offerings, often creating new business models and revenue streams.

- **Improves customer interactions with AI conversational tools.**
- **Drives autonomous logistics with robotics and vehicles.**
- **Delivers personalized recommendations to boost user experience.**
- **Expands digital workforce capabilities for more strategic focus.**

GenAI Categories



Enterprise AI

An organization implementing AI throughout the enterprise achieves transformative results and is seen as a lighthouse of innovation.



External AI

Externally facing AI solutions enable organizations to improve customer service, automate support tasks, and provide information in real time.



Pilot / Proof of Value AI

Evaluate AI use cases and relevant data to implement within business processes, aiming to demonstrate tangible value.



Internal AI

Implementing AI and automation capabilities to tackle internal use cases that empower users with data, enable proactive decision-making, and boost productivity.



Ethics

- Transparency & Exploitability
- Accountability
- Bias & Integrity



Organization

- Organizational Design
- Talent Model
- Culture



Process

- ROI & Funding
- AI Governance
- Delivery



Technology

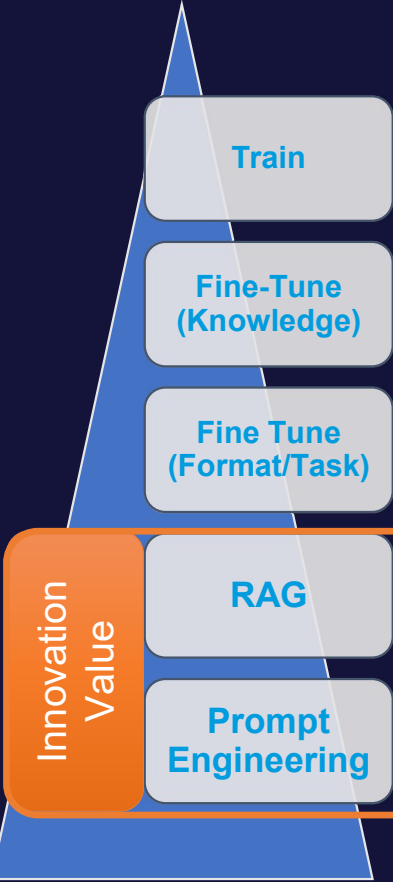
- Tools & Architecture
- Security & Continuity
- Deployment Models



Data

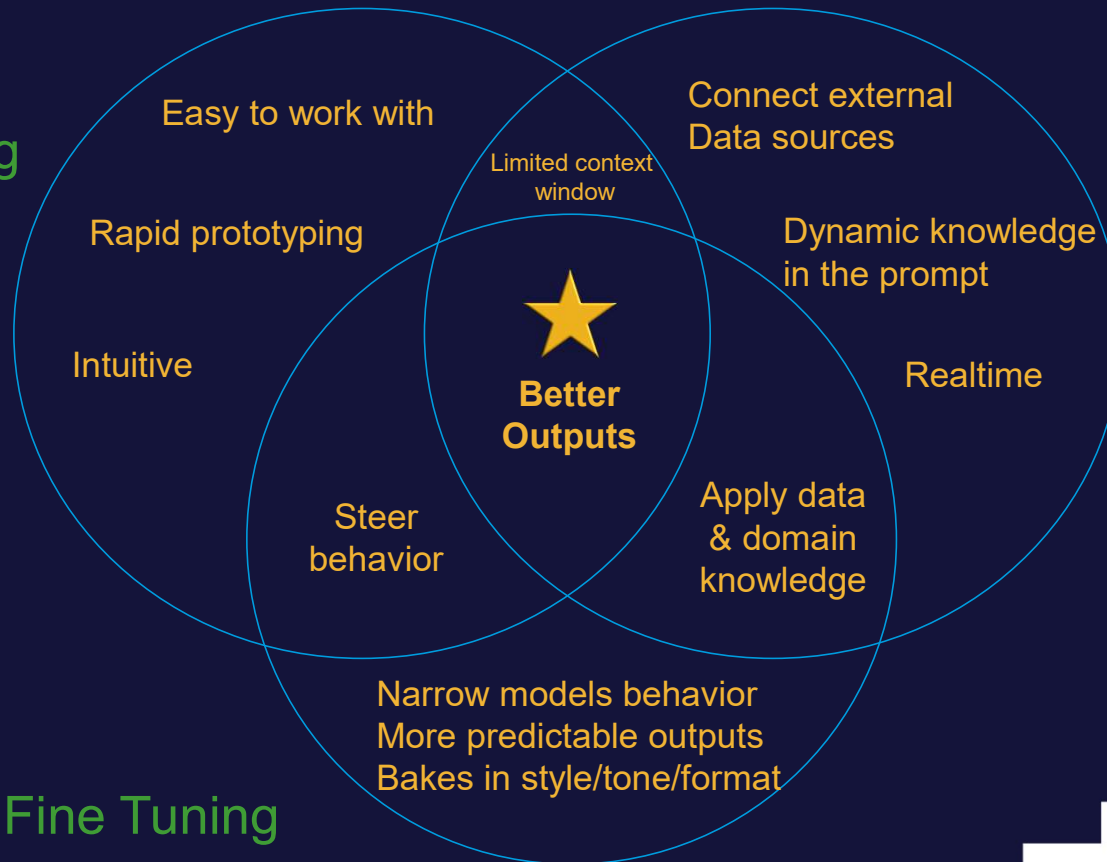
- Enterprise Data Strategy
- Data Management
- Data Governance

Bring your context to a Large Language Model

		Process	Cost	Data	Effort	Cost	Complexities	Time to market	Skillset Required
 <p>Innovation Value</p>	Train	Create a new model and train it using your company's data	Tens of millions	Trillions of tokens	Extremely High	Extremely High	Tailoring, high costs, many experts, data gathering, algorithms, tech setup.	Months to Years	Expert in ML, Data Engineering
	Fine-Tune (Knowledge)	Take a pre-trained model and fine-tune it using your company's data.	millions	Millions of tokens	Very High	Very High	Re-add safety, align data, add functions.	Weeks to Months	ML expertise, Data alignment
	Fine Tune (Format/Task)	Take a pre-trained model and fine-tune for a specific format or task.	Tens of thousands	100k's of tokens	Moderate	High	Ensure model-system fit, correct output, no new data.	Weeks	ML & Software Development
	RAG	Combine a LLM with a retrieval system.	Tens of thousands	Depends on your data	Moderate	Low	Keep data consistent across systems.	Weeks to Months	Data Engineering, Some ML
	Prompt Engineering	Craft specific prompts or input structures to guide the model's output	Cents	Length context	Extremely low	Low	Outputs may not always meet expectations.	Days to Weeks	Basic Coding, NLP Understanding

Methods to improve accuracy

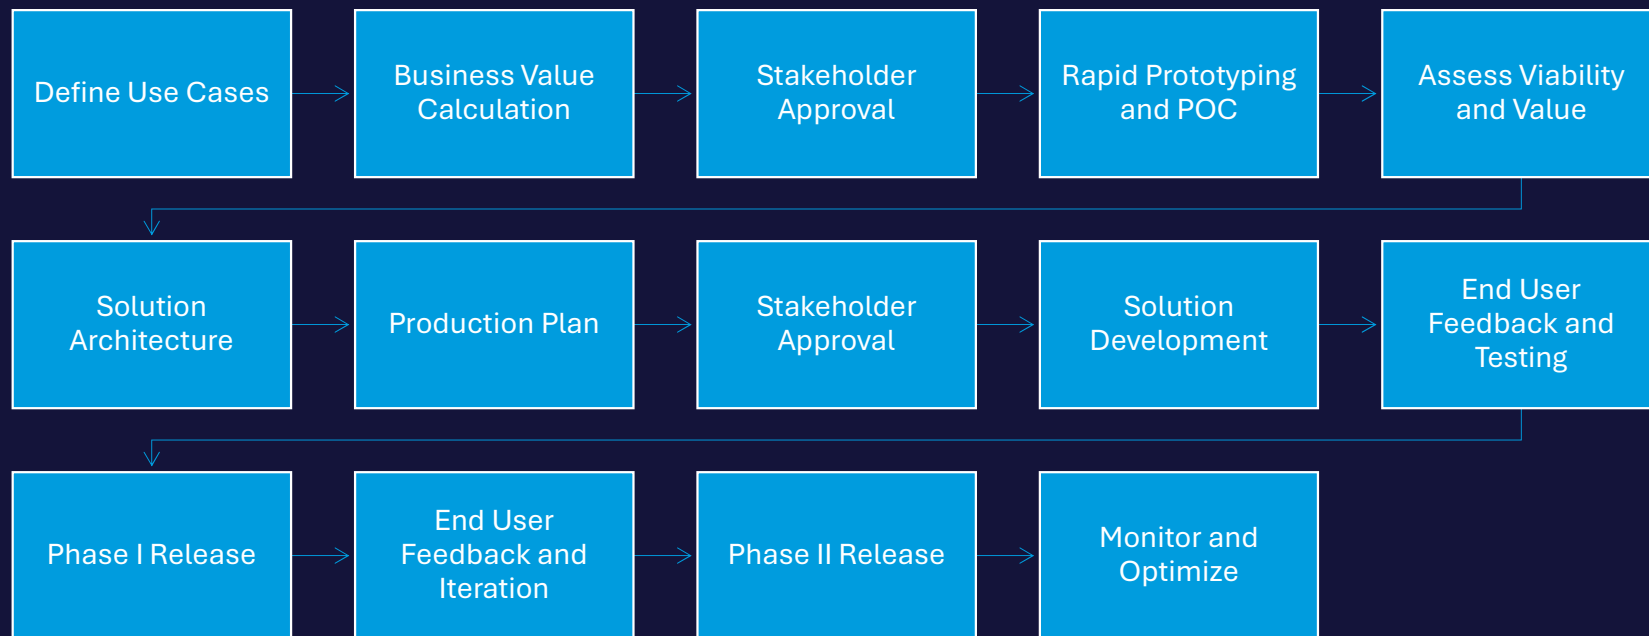
1. Prompt Engineering



3. Retrieval Augmented Generation (RAG)

2. Fine Tuning

Build an enterprise AI application



Risk dimensions in AI implementation

Understanding potential risks becomes pivotal as businesses adopt AI for innovation and operational efficiency. AI-related risks span across several dimensions, from technical and regulatory to ethical and reputational concerns. Acknowledging and planning for these risks can facilitate a smoother integration of AI, ensuring responsible and effective use while also maintaining public trust and compliance with emerging regulations.

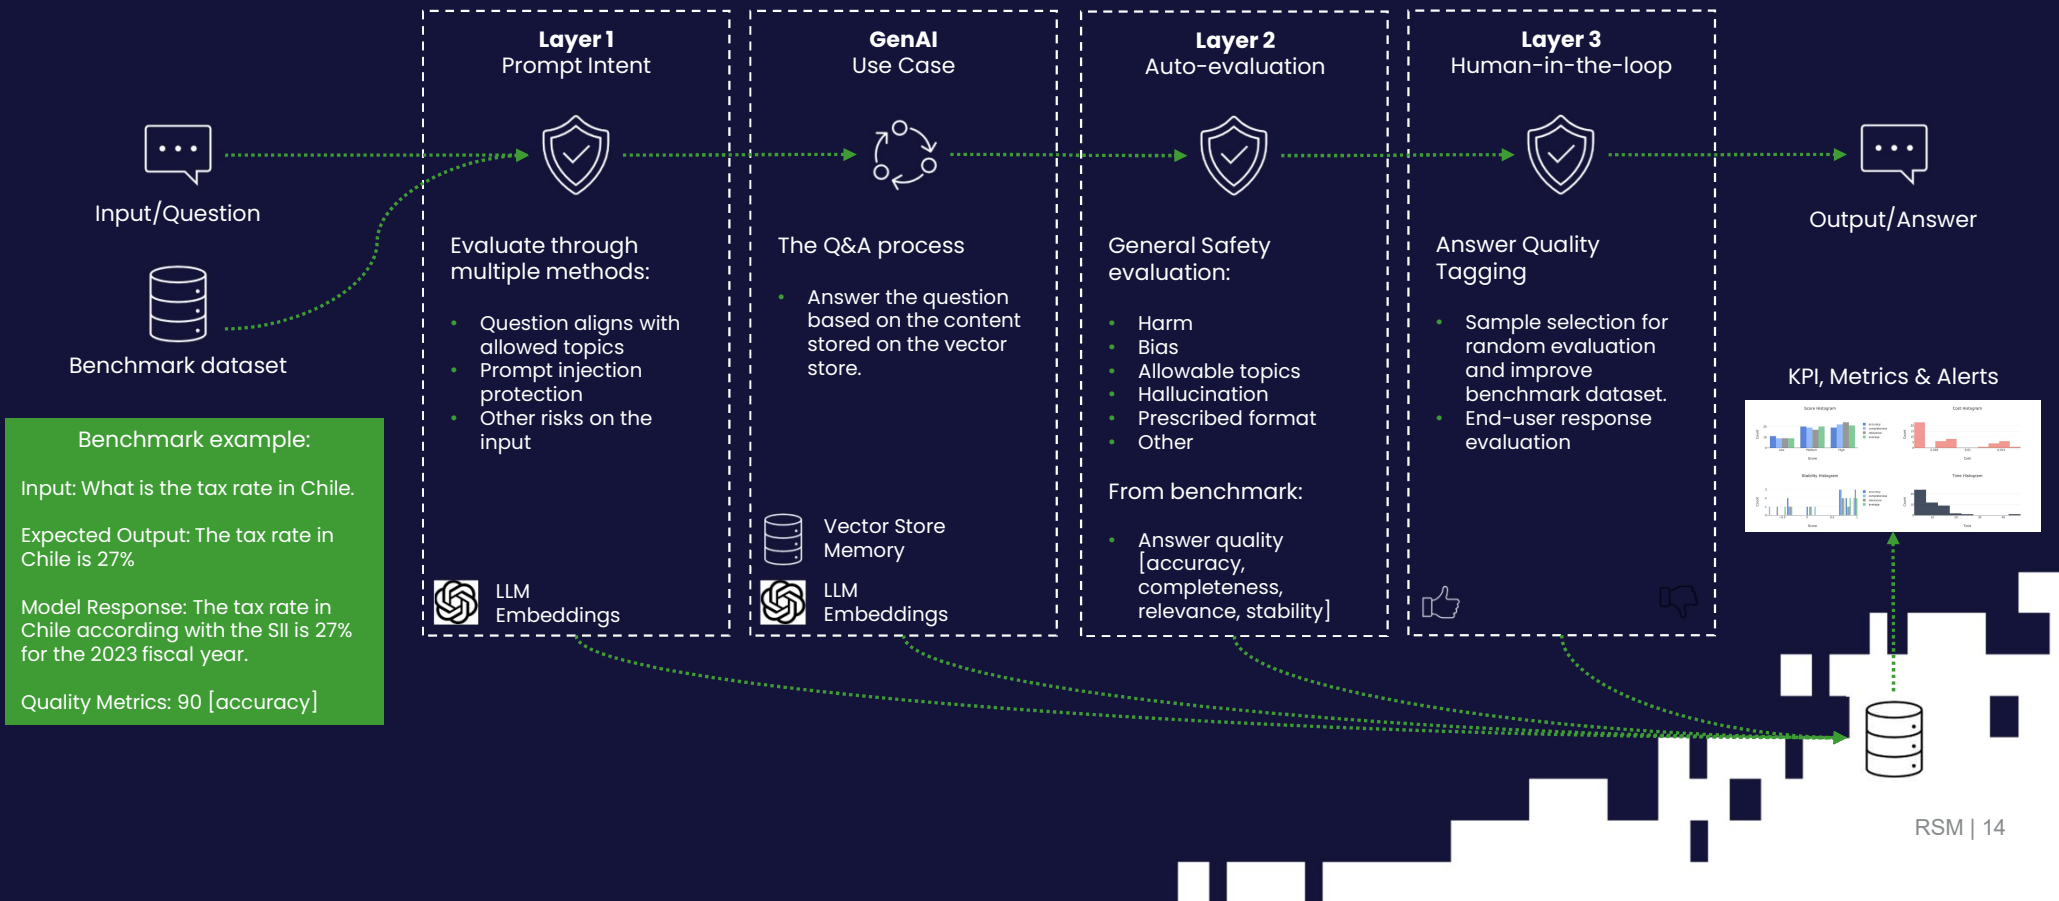
Technical & Semantic	Legal & Regulatory	Ethical	Reputational	Operational & Financial
<ul style="list-style-type: none"> • Inaccuracy: AI systems may produce incorrect or misleading results. • Bias: AI models can perpetuate biases in the training data, leading to unfair outcomes. • Security: There is a risk of data breaches or misuse of AI for malicious purposes. • Reliability: Dependence on AI can lead to risks if the system fails or performs poorly. • Obsolescence: Fast AI evolution can render systems obsolete, requiring expensive, regular upgrades. 	<ul style="list-style-type: none"> • Compliance: As regulatory frameworks for AI are emerging, non-compliance can lead to penalties and reputational damage. • Privacy: Misuse of AI could violate privacy laws or lead to unauthorized data sharing. • Liability: Determining responsibility for decisions made by AI can be complex, potentially leading to legal disputes. 	<ul style="list-style-type: none"> • Fairness: The use of AI could potentially exacerbate social inequalities if it is biased or misused. • Transparency: Lack of transparency in AI decision-making can lead to mistrust and ethical concerns. • Job displacement: AI-led automation could lead to job losses and social disruption if not managed responsibly. 	<ul style="list-style-type: none"> • Public Perception: AI misuse or failures damaging the company's image. • Trust: Potential loss of consumer trust due to irresponsible data handling or adverse AI decisions. • Responsibility: Unclear AI accountability leading to reputation damage. • Ethical Controversies: Public backlash from lapses in AI ethical practices. 	<ul style="list-style-type: none"> • System & Integration Issues: Failures and integration problems causing disruptions and cost overruns. • Data Management & Scalability: Poor data handling and scaling issues limiting progress and adding costs. • Costs & Regulatory Fines: Large initial and ongoing costs, coupled with potential regulatory fines.

Shape the AI threat vector analysis

Understanding threat vectors in AI implementation is pivotal to enhancing AI system security, resilience, and effectiveness through risk mitigation, defense strategy, compliance, and legal considerations.

Adversarial Attacks and Evasion	Data-Related Risks	Model Security and Intellectual Property	System-Level Risks	Human Capital and Talent Challenges
<ul style="list-style-type: none"> • Adversarial Attacks: These attacks exploit vulnerabilities in AI models by adding small, carefully crafted perturbations to input data. The goal is to deceive the model into making incorrect predictions. • Evasion Attacks: Aim to fool an AI system during inference. Attackers modify input data to evade detection or classification. • Inference Attacks: Target the data output rather than the model itself. By observing the responses, an adversary can infer sensitive information about the training data or the model's behavior. 	<ul style="list-style-type: none"> • Data poisoning: Manipulating training data to corrupt the AI model's output. • Exploiting Biases: AI systems can inherit biases from their training data. Exploiting these biases can lead to unfair or discriminatory outcomes, affecting certain groups disproportionately. • Dependence on External Data Sources: AI models often rely on external data (e.g., APIs, databases). The model's performance is affected if these sources fail or provide inaccurate data. 	<ul style="list-style-type: none"> • Model Stealing: In this attack, an adversary tries to replicate or steal a trained model by querying it and using the responses to create a similar model. It's a form of intellectual property theft. • Insider Threats: Malicious insiders (employees, contractors, etc.) can intentionally manipulate AI systems or leak sensitive information. 	<ul style="list-style-type: none"> • AI System Misuse: This risk involves using AI systems for unintended purposes, such as weaponizing or using them to violate privacy rights. • AI System Failure: AI models can fail due to unexpected inputs, data drift, or other factors. Understanding failure modes is crucial for robust deployment. • Regulatory and Compliance Risks: Organizations must comply with legal and ethical guidelines when deploying AI systems. Non-compliance can lead to legal consequences. 	<ul style="list-style-type: none"> • AI Talent Risks: The shortage of skilled AI professionals poses a risk. Organizations need to manage talent acquisition, retention, and skill gaps.

Gen AI – Defense in Depth



Governance Intersection

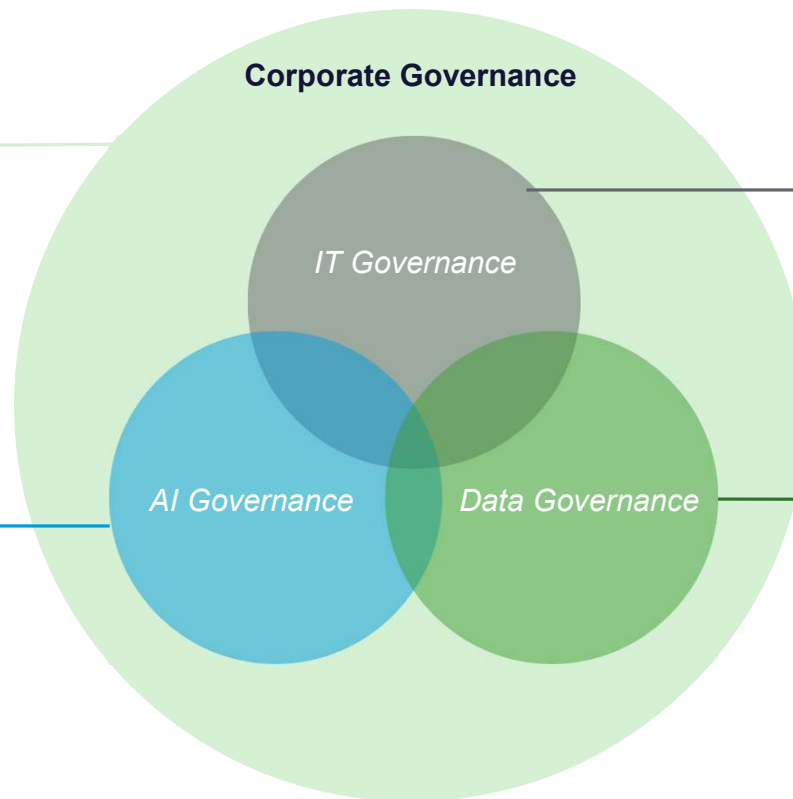
AI Governance is linked to and overlaps with IT and Data Governance. The Enterprise AI and Governance Strategy that is developed should established structures, processes, and procedures within these realms.

Corporate governance refers to the rules, practices, and processes by which a company is directed, controlled, and managed. It ensures the company is run responsibly, ethically, and sustainably.

Corporate governance helps to **build trust and confidence among stakeholders**, **reduces the risk** of corporate scandals and fraud, and **enhances the company's reputation** and ability to attract and retain investors.

AI governance ensures that AI is developed and used responsibly and ethically, is aligned with societal values and serves the public interest.

AI Governance involves establishing clear **accountability for AI decisions**, ensuring that AI is **transparent and explainable**, and **addressing bias, fairness, and discrimination** in AI systems.



IT governance refers to the processes, policies, and frameworks organizations implement to ensure that their information technology (IT) functions effectively and supports their business objectives.

IT governance helps organizations **make better use of their IT investments**, reduce risks associated with IT, and **improve the overall efficiency and effectiveness** of their operations.

Data governance ensures accuracy, completeness, consistency, availability, and security, enabling the organization to derive maximum value.

Data governance involves establishing data standards, defining data quality rules, and developing data collection, storage, processing, and sharing processes. It also includes **assigning roles and responsibilities for data management**, establishing data ownership and stewardship, and **ensuring compliance** with regulatory requirements related to data privacy and security.

RSM AI governance

Artificial intelligence solutions have immense potential to improve your business, but risks can derail progress

RSM’s experienced AI risk professionals provide comprehensive solutions to help you establish an effective AI governance program before implementing technology. AI governance goes beyond typical corporate, data, or IT governance by considering additional generative AI factors. AI solutions, especially generative AI, require guardrails and controls to ensure quality output. It’s crucial that AI processes align with your business strategies and regulatory guidelines, and our professionals are here to ensure that.

RSM’s proven AI Governance Framework guides our clients’ risk-based approach. It constantly evolves and integrates critical elements from many best-practice frameworks to create a more comprehensive approach covering more potential risks than a single framework. While extensive, the framework is also adaptable, with the level of governance tailored to align with the risk you are exposed to.



Security Layers for Generative AI

Cross Company: AI Policy | Model Risk Management | Employee Communication | Culture of Accountability | Feedback Channels

	End User	Product	Provider	GenAI Model
	End users utilizing GenAI products to interact with audit tools and models.	The GenAI-powered application, like Copilot or ChatGPT, is designed to assist or automate tasks.	The company is responsible for developing and providing access to a GenAI model, such as OpenAI, Microsoft Azure, or AWS.	GenAI employs large language models (LLMs) like GPT, Gemini, and Claude to comprehend and generate human-like text, enabling intelligent responses.
Risks	<ul style="list-style-type: none"> Misinterpretation of results Excessive dependence on model output Ineffective prompt crafting Manipulation of prompts Exposure of sensitive information. 	<ul style="list-style-type: none"> Exposure of sensitive data Bugs, defects or malfunctions Misconfiguration Unexplainable results 	<ul style="list-style-type: none"> Data breach Corporate data leakage into training data Model unavailable due to GPU shortage 	<ul style="list-style-type: none"> Limited training data / obsolete Unpredictable output Inaccurate & Biased outputs Model performance drift
Safeguards	<ul style="list-style-type: none"> User education / training Application input controls Output verification Data loss prevention alerts “Acceptable Use” policy Data handling protocols 	<ul style="list-style-type: none"> DevSecOps practices Data encryption, tokenization/data masking User access & entitlements Periodic security reviews Change control Application monitoring 	<ul style="list-style-type: none"> Third-party Risk Assessment License agreement, terms and conditions SOC report Service Level Agreements 	<ul style="list-style-type: none"> Performance testing against the standard test dataset Development documentation and release notes Red Team pressure testing User feedback

AI governance organization structure

Role	Assignment	Responsibilities	Members
AI Governance Board	Strategic oversight and policy-making	Approves AI strategy, policies, and major projects. Ensures alignment with business objectives and compliance with legal and ethical standards.	C-Level Executives, Head of AI, Legal Advisor, Ethics Officer.
AI Project Management Office (PMO)	Coordination and management	Oversees project planning, resource allocation, and execution. i.e., Projects align with organizational strategy and governance policies.	AI Project Managers, Technical Leads
AI Ethics Committee	Guidance on ethical considerations	Review projects for ethical implications, provide recommendations, and guide adherence to ethical standards.	Ethics Officer, Legal Advisor, External Ethicists, AI Developers.
Data Governance Team	Management of data assets	Ensures data quality, security, and privacy. Develops and enforces data governance policies.	Data Governance Lead, Data Scientists, Data Protection Officer, IT/Security.
IT and Cybersecurity Department	Security operations and technical support	Provides IT infrastructure, implements cybersecurity measures for AI systems, and monitors for security threats.	CIO, IT Personnel, Cybersecurity Experts.
AI Research and Development (R&D) Team	Solutions Development	Conducts AI research, develops and tests AI models.	AI Engineers, Data Scientists, Research Specialists.
Compliance and Legal Department	Legal oversight	Ensures AI initiatives comply with laws and regulations, manages legal risks.	Legal Advisor, Compliance Officers.
Human Resources (HR)	Workforce management	Manages AI-related training and development, addresses workforce implications of AI deployment.	HR Managers, Training Coordinators.

Next steps



Is your organization getting the benefits of AI solutions?

Whether you prefer to learn at your own pace, use interactive tools for a hands-on learning experience, or have a conversation with someone who can guide you, we have you covered!



Get a free assessment using AI

Use our AI-powered guide: The **AI & Digital Maturity Compass**. Gain personalized insights on using AI to cut costs, boost sales, and improve decision-making.

 [Get a personalized report on using AI](#)



Explore AI trends and resources


Explore the latest AI adoption trends, case studies, frameworks, videos, and implementation guides for successful and responsible use of AI in your organization.

 [Explore AI content and tools](#)



Have a no cost consultation

After you explore our content and try our AI and Digital Maturity Compass, contact us for a no-cost consultation with a Microsoft expert.

 [Start a consultation with an expert](#)

Protect data and privacy

AI needs access to data to deliver value. Use reputable providers.

When you use [Copilot for M365](#), [Copilot Studio](#) or [Azure AI Studio](#), your data is never used to train AI models, and it remains confidential and safe.

[Learn more here](#)



1

Sign up for the [no cost AI Consultation with RSM](#)

Select your preferred method to contact on question #12, then select **RSM** on question #15.

2

Utilize the no cost [AI & Digital Maturity Compass assessment](#) to get a personalized report on how AI can help you and your business today.



3

Access the free [AI for Small to Midsize Businesses webpage](#) to learn more about AI resources and tools to help your business.

Thank you





THE POWER OF BEING UNDERSTOOD ASSURANCE | TAX | CONSULTING

This document contains general information, may be based on authorities that are subject to change, and is not a substitute for professional advice or services. This document does not constitute assurance, tax, consulting, business, financial, investment, legal or other professional advice, and you should consult a qualified professional advisor before taking any action based on the information herein. RSM US LLP, its affiliates and related entities are not responsible for any loss resulting from or relating to reliance on this document by any person. Internal Revenue Service rules require us to inform you that this communication may be deemed a solicitation to provide tax services. This communication is being sent to individuals who have subscribed to receive it or who we believe would have an interest in the topics discussed.

RSM US LLP is a limited liability partnership and the U.S. member firm of RSM International, a global network of independent assurance, tax and consulting firms. The member firms of RSM International collaborate to provide services to global clients, but are separate and distinct legal entities that cannot obligate each other. Each member firm is responsible only for its own acts and omissions, and not those of any other party. Visit rsmus.com/aboutus for more information regarding RSM US LLP and RSM International.

RSM, the RSM logo and *the power of being understood* are registered trademarks of RSM International Association.

© 2023 RSM US LLP. All Rights Reserved.