

A lighthouse on a rocky shore at sunset, with a beam of light shining across the water. The sky is a mix of orange, red, and blue, and the water is dark blue with a reflection of the light.

AI Governance  
**Lighthouse  
Case Study**



KPMG Australia

# Key Governance Insights



The responsible AI guardrails and controls that an organisation put in place need to be specific to the system and context for deployment, including the unique needs and interests of stakeholders, and the particular risks associated with each use case. Careful experimentation with, and gradual deployment of AI systems, provide opportunities for an organisation to develop and identify the necessary guardrails and controls.



High quality data is critical for achieving the best results with Generative AI. Organisations will need to consider uplifting their existing data governance processes so that ownership rights, lineage, provenance, and bias are assessed. This data governance process is a key factor to ensuring system outputs are reliable, explainable, and lawful.



The development, deployment and governance of AI systems is more likely to succeed where an organisation draws on the skills and multidisciplinary expertise of its people (bottom-up), is supported by an engaged leadership team (top-down), and is subject to ongoing monitoring and improvement.

# Essential components of AI Governance covered

1. Accountability, oversight, roles & responsibility



2. Governance structures



3. People, skills, values and culture



4. Principles and policies



5. Practices, processes and controls



6. Supporting infrastructure



7. Stakeholder engagement, co-design and impact assessment



8. Monitoring, reporting and evaluation



# How is KPMG Australia using AI?





<b>Firm</b>	KPMG Australia
<b>Employees</b>	Over 10,000
<b>Revenue</b>	\$2.553 bn in FY 2023
<b>Sector</b>	Professional services

KPMG Australia (KPMG) sees a significant opportunity for AI to support its employees, create value for clients, and deliver increased innovation. A leading example of this is its development and use of its internal digital assistant and proprietary version of ChatGPT, KymChat.

In March 2023, KPMG launched KymChat for internal use for its employees. KymChat was originally designed to assist staff in finding the right expert across the business. Based on its success and flexibility, its functionality was expanded. KymChat is now KPMG's most widely used AI assistant and is used daily by a significant of the firm's 10,000 staff for a variety of uses, including answers questions about internal policies, compiling thought leadership, and in approved cases – preparing draft tax advice to clients.


KymChat attracted much interest from member firms around the globe demonstrating the Australian firm's capabilities as a leader in AI. Microsoft invited KPMG to showcase KymChat at the Microsoft 2023 Ignite conference in Seattle, attended by over 25,000 industry and technology experts.

# How did KPMG develop KymChat?





As a professional services firm, security and data governance are core to KPMG's business model and culture.



Paramount to KPMG is the importance of protecting client and firm data, so initially staff access to the web version of ChatGPT was blocked due to significant concerns that data might end up shared on ChatGPT's public platform (given OpenAI can retain and use data entered into the public version of ChatGPT). At the same time, KPMG recognised the value of generative AI and a growing demand from employees to make use of the technology. As Martin Wing, AI Strategic Lead, commented, 'there was a need for us to explore the potential of AI in a safe and secure way...to explore the art of the possible'.



## Following some controlled experimentation with ChatGPT KPMG decided to develop its own chatbot, KymChat.

KPMG partnered with Microsoft to build an internal facing generative AI system on Microsoft Azure's OpenAI Service, allowing the firm's employees to explore this modern technology without client data leaving the KPMG environment and maintaining KPMG standards for privacy, confidentiality, and data protection.

KPMG's vision was to create an AI agent that would enhance the efficiency and effectiveness of its employees. From the beginning, KymChat was not just a technology led initiative, but business led as well. As an initial, low-risk, step KymChat was carefully developed by a cross-functional KPMG team that drew from the expertise of the KPMG's technology, privacy, risk, legal, data, and product teams. As Robert Finlayson, KymChat Product Manager commented, 'from the get-go, we knew the guardrails we needed as we had everyone in the same room'.

As part of KPMG's commitment to continuous improvement, KymChat is tested, monitored, and refined on an ongoing basis to ensure it is delivering reliable, safe, and helpful responses. Whilst originally created to focus on identifying internal experts, KPMG has designed KymChat's different functionality around different 'personas', indicating a different area of expertise for a deployed and trained model.

KPMG's phased approach to the build out of KymChat has allowed the firm to evolve its AI governance structures and processes in step with the increased sophistication of its AI tools in a way that meets the increasing opportunity, complexity, and risks associated with AI.



# What are the KPMG personas?





KymChat has several different personas in operation and more in development. The functionality of each persona is based on its access to different data sources and instructions on how it interacts with prompts.

Careful consideration is given to the data being used by that persona, the purpose of that persona, and which staff are given access, providing effective data governance controls.



**Integra:** An enterprise-wide persona that assists KPMG staff in finding relevant policies, procedures, and guidelines, identifying firm leaders and subject matter experts, and the summarization of the collective firm's published thought leadership.



**Authentica:** A safe and secure version of ChatGPT 4.0 where any information entered remains within the KPMG environment.



**Sustania:** An Environmental, Social and Governance (ESG) persona that draws on KPMG's thought leadership on ESG issues to answer questions and draft new content.



**Ignition:** Draft proposals and sales documents helping the sales teams with new client and business opportunities.

# How is KPMG approaching AI governance for KymChat?





## KPMG has taken a multi-layered and sequenced approach to AI governance for KymChat.

The initial KymChat was a low-risk foray into generative AI, subject to the firm's existing standards, frameworks and principles around data, cyber security, privacy and other controls. Over time and keeping pace with the evolution of KymChat, the firm has begun to develop a robust, AI-specific governance structure that allows the firm to confidently explore the opportunities and address the risks that may emerge with the scaling of KymChat and other AI tools. At the core of the firm's approach is the KPMG Trusted AI framework which addresses potential risks to people and the planet and the operationalisation of Trusted AI principles through guardrails and controls. Important guardrails surrounding KymChat are explored below.

- » Careful use of data
- » Policies and disclaimers
- » Capability building and user-focused design
- » Impact assessment
- » Monitoring, reporting and evaluation



## KymChat guardrails

» Careful use of data

» Policies and disclaimers

» Capability building and user-focused design

» Impact assessment

» Monitoring, reporting and evaluation

### » Careful use of data

KPMG worked with alliance partner Microsoft to gain insights and knowledge from their Azure OpenAI studio resources and technical support, obtain expert system architecture advice, and learn top industry practices. This collaboration ensured that data storage remained under KPMG's secure control, with internal data access strictly regulated according to employee work relevance.

KPMG follows strict data governance processes to assess data for ownership rights, lineage, provenance and bias to ensure they are not inappropriately or unlawfully using anyone's data. The following steps are taken:

- **Permission:** KPMG does not include data within KymChat without investigating and confirming that it has the right and permission to use that data.
- **Anonymisation:** If KPMG materials prepared for a client are to be included into a specific KymChat persona, in accordance with KPMG Terms and Conditions and where consent is provided, those materials are sanitised to ensure that any identifying or confidential information is removed.
- **Legal review:** KPMG's legal team provides sign off on the use of data by KymChat.



## KymChat guardrails

» Careful use of data

» Policies and disclaimers

» Capability building and user-focused design

» Impact assessment

» Monitoring, reporting and evaluation

### » Policies and disclaimers

The use of KymChat by KPMG employees is governed by many internal policies, including terms of usage, acceptable use, privacy and data. To ensure that staff properly understand these risks, their obligations, and the limitations of KymChat, clear disclaimers are included in the application.

### » Capability building and user-focused design

Beyond the technical implementation, one of the goals of the KymChat team is to ensure high levels of adoption and retain its user base. To facilitate this, KPMG has provided training to all its Australian staff on KymChat, what it can do, and how to engineer successful prompts. KymChat also provides users with suggested prompts so they can see clear examples of how to structure prompts and the answers these deliver (including common questions).



## KymChat guardrails

» Careful use of data

» Policies and disclaimers

» Capability building and user-focused design

» Impact assessment

» Monitoring, reporting and evaluation

### » Impact assessment

KPMG assessed the impact of KymChat through the [Microsoft Responsible AI Impact Assessment Template](#) combined with KPMG's own Trusted AI principles. By applying this assessment process, it introduced additional features to mitigate potential harms such as improvements to user experience, increased transparency and greater alignment to business process. Through the value demonstrated in applying the AI Impact Assessment to KymChat, KPMG's internal risk team are now working on embedding it into the firm's formal risk management processes.

### » Monitoring, reporting and evaluation

KPMG emphasised the importance of monitoring and testing. The responses from generative AI will not always be the same and carry the risk of providing false information or hallucinations. High levels of accuracy are important for KymChat users, as incorrect responses increase the review time and erode any potential productivity benefits. Accordingly, it is very important to monitor, test and fine tune.

A key metric that KPMG uses to test KymChat's performance is the KPMG Prompt Confidence Index. This metric involves measuring KymChat's response against the correct response and giving it a score for how well it performed against a series of different attributes. KPMG has tested prompts against hundreds of questions. It has undertaken work to improve the scores, such as improving the precision, increasing the data sources, and updating the interaction with the Large Language Model. KymChat is now answering with over 90% confidence.

# What can we learn from KPMG's use of AI?







## What can we learn from KPMG's use of AI?

1 The value of careful design

2 Importance of data governance

3 Leadership

### 1 The value of careful design

Although generative AI systems have risks and shortcomings, this should not prevent innovation. These shortcomings can be overcome by engaging in careful and thoughtful development, and the implementation of effective guardrails and governance specific to the particular risks posed by the system in its deployment context. KymChat was built for a clear purpose with KPMG's existing legal, risk and compliance requirements front of mind and has evolved in concert with KPMG's maturing AI governance approach. It was designed by professionals with expertise in these areas so that necessary safeguards were built into the system. In comparison, many organisations are rushing to adopt AI systems without cautious consideration. This exposes organisations to potential commercial, legal and reputational risks. Organisations should first engage in careful experimentation with AI and ensure that the risks of technology are appropriately considered before widespread adoption and deployment.



## What can we learn from KPMG's use of AI?

1 The value of careful design

2 Importance of data governance

3 Leadership

### 2 Importance of data governance

KPMG's experience shows that it is hard to overstate how important data governance is for generative AI. The temptation is often to develop these systems by giving them as much data as possible. However, much of an organisation's existing internal knowledge, such as policies, precedents and other work products, is unstructured and has never formally had data governance applied to it. As such, it may be the first time that issues like data quality, provenance, lineage, rights and bias are considered. Without quality data governance, AI is unlikely to deliver effective and safe outcomes. Organisations need to ask themselves about the source of that data, whether it is a trusted source, the quality of the data and how to assess that, and whether the organisation has the right to use the data in its AI systems.

### 3 Leadership

Active support, engagement and direction from senior leadership in an organisation is needed when developing and deploying new and innovative AI solutions, particularly where they may substantially impact the work of employees and the business. Within KPMG, there is a clear and determined leadership that has championed the opportunity of generative AI and marshalled the firm's resources accordingly. The National Executive Committee is leading the firm's drive for innovation and is actively involved in helping its staff reinvent KPMG.

# What's next for KPMG?

As KPMG Chief Digital Officer John Munnely said, 'KymChat is just the start of KPMG's AI journey, and in time will be one of many products within our broader AI strategy.' KPMG globally is investing over \$2bn USD in AI and cloud services during the next five years as it expands its partnership with Microsoft.

KPMG is considering and developing further KymChat personas for specific teams and use cases. Whilst KymChat was initially developed by KPMG Australia, it has now been internally open sourced to all KPMG global firms. KPMG is also helping clients launch their own version of KymChat through KPMG's KymChat Accelerator solution, which combines the product with KPMG consulting services.



## Human Technology Institute

For more information on HTI's AI Corporate Governance Program,  
or to join our AI Governance network, please contact:

Llewellyn Spink, AI Corporate Governance Specialist  
[llewellyn.spink@uts.edu.au](mailto:llewellyn.spink@uts.edu.au)