KPMG

We pressure-tested our AI governance against a global standard. Here's what we found

If you can't prove your AI governance works, you're just hoping for trust.



kpmg.com

Introduction

AI is rewriting governance rules at a pace that makes old frameworks feel antique before they're even adopted. If generative AI was yesterday's headline and agentic AI is today's buzz, tomorrow's disruption is anyone's guess. In this environment, governance must sprint just to keep pace, and it has to work in practice, not just on a whiteboard.

The risky truth: Too many technology leaders think governance exists because they have a policy in place. But AI evolves so quickly that the ink is barely dry before that policy becomes obsolete.

At KPMG, we know we can't ask clients to adopt AI governance rigor without applying it ourselves. We operate across diverse industries, each facing its own AI risks and opportunities, and the credibility of our advice depends

The risky truth:
Too many technology leaders think governance exists because they have a policy in place. But AI evolves so quickly that the ink is barely dry before that policy becomes obsolete.

In short, we needed to prove to ourselves and to the market that our governance can withstand both the speed and scrutiny of modern AI. That belief led us to ISO 42001 certification.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044178-1A

ISO: What it is and why it matters

The International Organization for Standardization (ISO) is a global, nongovernmental body that develops and publishes international standards, including AI governance. ISO comprises representatives from more than 170 national standards bodies and its guidance sets the standard for quality and safety.

ISO 42001 is the world's first international standard for AI management systems.[1] It presents a framework for designing, implementing, and continually improving AI governance inside organizations. In plain language: It's the global benchmark for demonstrating your organization's management of AI systems is rigorous and effective.

For regulators, ISO 42001 offers a yardstick for compliance and oversight. For customers, it's tangible assurance that AI is being managed responsibly. For the market, it's a differentiator: Certification signals maturity, accountability, and readiness to scale AI safely.

That's why KPMG chose to pressure-test our own AI governance against ISO 42001. We didn't pursue certification for its own sake—we did it to lead by example for clients, to validate our governance against tough criteria, and to understand firsthand what it takes to walk the "responsible AI" talk.

The result: We became among the first Big Four US professional services firms to earn ISO 42001 certification.[2] And along the way, we learned lessons that apply to any organization navigating AI governance.

We didn't pursue certification for its own sake—we did it to lead by example for clients, to validate our governance against the toughest criteria available, and to understand firsthand what it really takes to walk the "responsible AI" talk.

The result: We became the first Big Four US professional services firm to earn ISO 42001 certification.

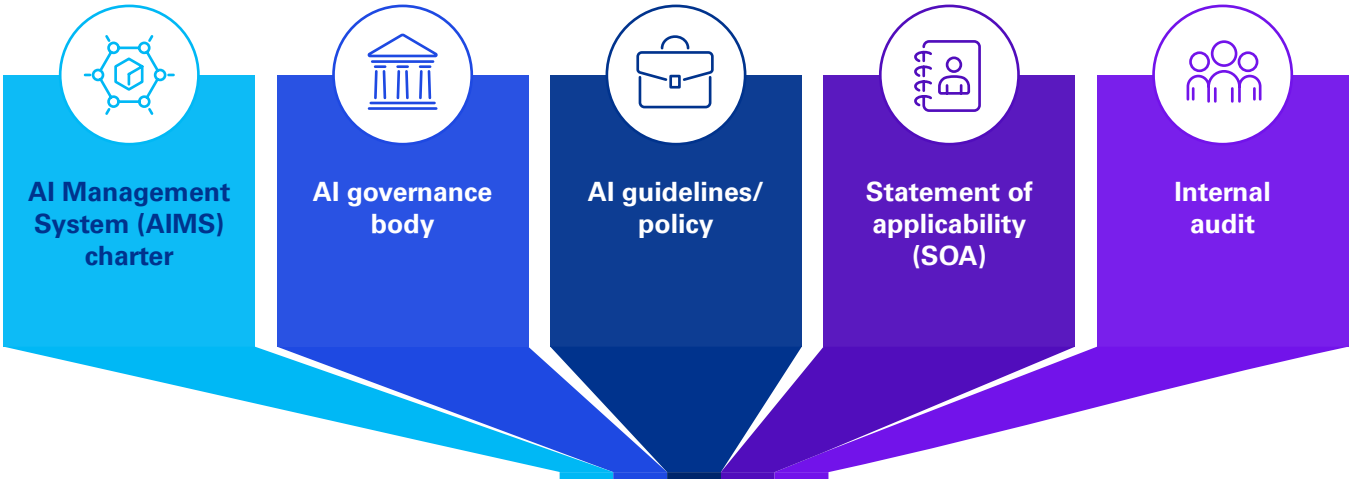
[1] "ISO/IEC 42001:2023," ISO.org, 2023.

[2] "KPMG LLP Among First of the Big Four in the U.S. to Receive ISO 42001 Certification," KPMG.com, November 19, 2025.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044178-1A

Certification pre-requisites



AIMS charter

Proprietary document that serves as an anchor for ISO 42001 by outlining the scope, function, and relevant interested parties of the US AIMS.

AI governance body

A designated group responsible for overseeing the AIMS and serve as "Management" within the context of ISO 42001. This body provides leadership, escalation paths, and accountability.

AI guidelines/policy

Set of principles, requirements, and expectations governing the development, deployment, and use of AI systems.

SOA

Document that lists all ISO 42001 Annex A controls, indicates which are applicable or excluded, and provides justification for each decision. The SOA serves as the response document and the authoritative control baseline for the audit.

Internal audit

Dedicated, separate function that conducts a review of the AIMS per the ISO 42001 standard before the certification audit. Internal audit validates that policies, controls, governance processes, and documentation are in place and operating effectively, identifying gaps that require remediation.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044178-1A

The ISO 42001 certification journey

ISO 42001 certification is a rigorous, multistage audit that probes the design and operational execution of AI governance. Here's how we navigated the gauntlet.

Step 1 | Assembling the task force

Our first move was assembling a cross functional certification task force. AI governance touches every corner of KPMG, so our team included specialists from Legal, IT, Risk, Compliance, and core business operations. The goal: Ensure governance is understood, adopted, and applied firmwide.

We started with a pre-readiness assessment aligned to 62 ISO standards and controls. This quickly revealed the scale of the challenge: When AI is infused into everything you do, governance applies to everything you build, use, and rely on.

At this stage, questions surfaced: How exactly had AI systems propagated through our everyday operations? Had oversight kept pace? Would our governance hold up under an external spotlight? We were already two years into a strategic AI program, but certification meant putting every assumption on trial.

Step 2 | Designing the framework and the Stage 1 audit

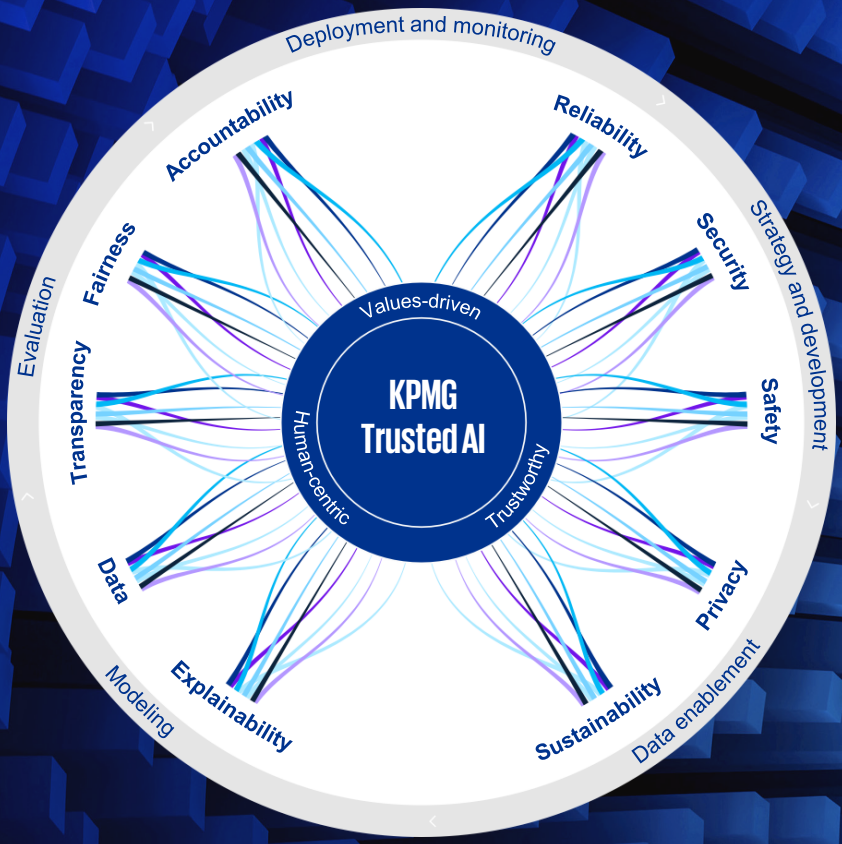
Stage 1 of the ISO audit examines the design of AI governance frameworks. Auditors look for policies, controls, and processes that are clearly documented and aligned with the standard. It's prescriptive enough to ensure rigor, but flexible enough that businesses can tailor controls to their objectives and risk profiles.

For KPMG, existing enterprise technology and risk management policies gave us a strong foundation, but ISO 42001 required us to connect those components into a more explicit AI management system. That meant aligning our enterprise controls with the firm's Trusted AI Guidelines, which are built around 10 pillars for responsible design, deployment, and use of AI. The result was a governance framework that could be consistently applied across the AI lifecycle, while still allowing enough flexibility to adapt as AI capabilities, risks, and regulatory expectations continue to evolve.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044178-1A

The Trusted AI framework



Our guiding philosophy: Governance should enable and accelerate, not choke, innovation. By embedding Trusted AI pillars from the start, we could advance trust and innovation in parallel. But we also learned that governance models that we thought were mature can be unevenly implemented and lack necessary documentation. And even more importantly, policies don't create maturity unless they're socialized consistently across the enterprise.

The KPMG Trusted AI framework is built around 10 core pillars: Accountability, Data, Explainability, Fairness, Privacy, Reliability, Safety, Security, Sustainability, and Transparency. These pillars help ensure AI systems are designed and operated responsibly, emphasizing ethical oversight, robust data governance, clear logic and outcomes, nondiscrimination, regulatory compliance, consistent performance, protection from harm, resilience against threats, environmental responsibility, and openness throughout the AI lifecycle. Together, they provide a thorough approach to developing, deploying, and managing AI that is trustworthy, human centric, and aligned with organizational values and regulatory expectations.

KPMG

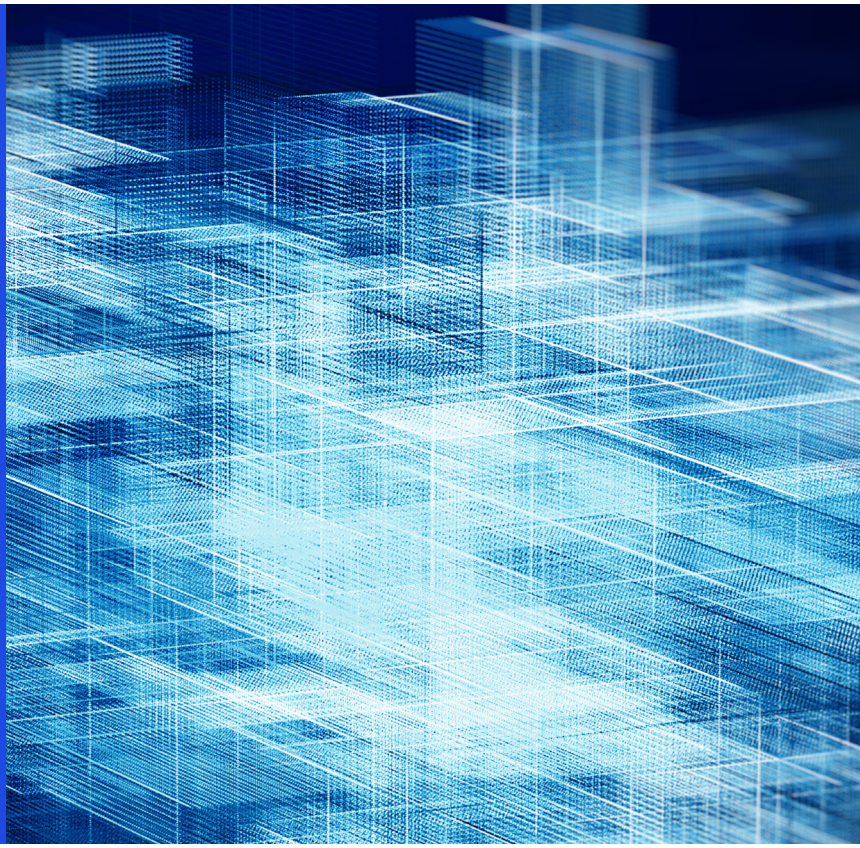
© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044178-1A

Step 3 | Proving it works during the Stage 2 audit

Stage 2 moves from intent to execution. Auditors evaluate whether governance is operationalized in real AI systems. This is where some organizations stumble: Passing Stage 1 means you have a governance model, but passing Stage 2 means you live by it.

We demonstrated governance in action through:

- An AI inventory
- System design documentation
- Risk evaluations and treatment
- Ongoing monitoring protocols
- Control group oversight



The test reinforced one of our internal strengths: a natural bias toward "doing the right thing" and high expectations for governance. But it also highlighted opportunities, including tethering disparate practices for efficiency and further documenting evolving governance standards for consistency.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044178-1A

An example of how the ISO 42001 certification effort drove maturity in our AI management system is the formalization of a quarterly AI system monitoring process. This structured approach has broad impact by:

Formalizing reassessment and documentation refresh cycles

Operationally, this introduces a disciplined cadence where each AI system is periodically revisited against its original purpose, risk classification, and performance expectations. System owners are prompted to validate that documentation—such as system descriptions, data sources, and risk assessments—remains current, creating a living record rather than a point-in-time artifact.

Ensuring changes are flagged and reflected in inventory records

As systems evolve—through model updates, new integrations, or expanded use cases—these changes are systematically captured and reconciled within the enterprise AI inventory. In practice, this results in improved traceability, where stakeholders can quickly understand what has changed, why it changed, and how it may impact risk or compliance obligations.

Integrating key functional groups into AI system decommissioning workflows

The process also strengthens cross-functional coordination when AI systems are retired or replaced. Legal, risk, technology, and data stakeholders are engaged through defined workflows to confirm that decommissioning is executed in a controlled and compliant manner—addressing data retention, access revocation, and downstream dependencies.

What this looks and feels like operationally

This level of structure shifts AI governance from a reactive exercise to a repeatable, embedded business process. System owners have clear expectations and checkpoints, governance teams gain greater visibility and consistency, and leadership benefits from more reliable, decision-ready insights into the AI portfolio. Ultimately, governance becomes less about periodic review and more about continuous, operationalized oversight that scales with the organization's AI footprint.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044178-1A

Beyond certification: continuous improvement and third party accountability

Certification isn't the finish line. It is an external validation of our governance foundation and a catalyst for continued improvement. We are expanding policies to more explicitly address AI components, increasing awareness through targeted training, and strengthening resilience against emerging AI risks.

One critical lesson: AI governance doesn't stop at what you build in house. Organizations increasingly rely on third party models, platforms, and software as-a-service solutions, but responsibility doesn't outsource alongside the technology. We remain accountable for outcomes and risks from AI systems that incorporate third party components, and so should you.

Practically, that means applying due diligence and ongoing oversight to suppliers and partners. Trusted AI governance must extend across your value chain because clients, regulators, and the public don't care if the source was internal or vendor built—they care whether it's safe, reliable, and responsible.

Is ISO 42001 a one-time certification, or is there ongoing work required?

Rather than a one-time achievement, ISO 42001 is a commitment to continuous improvement. Achieving the certification sets off a three-year certification period, where the controls are evaluated annually against a sampling of in-scope systems. It also includes an annual review and refresh of the AI management system standards and controls responses that accounts for maturation and evolution of the enterprise. Organizations must regularly review and enhance their AI governance framework to adapt to new technologies, evolving regulations, and lessons learned from their own AI deployments. The standard is designed to foster a culture of responsible and accountable AI management that matures over time.

Five lessons for your enterprise

1. **External benchmarks reveal blind spots**
   Even mature organizations discover gaps when governance is tested against global standards

2. **Cross functional teams are nonnegotiable**
   AI governance impacts many disciplines; representation helps ensure comprehensive coverage.

3. **Governance maturity is social, not just structural**
   Policies should be understood and embraced across the organization

4. **Third party AI requires the same governance as in house builds**
   Accountability extends beyond ownership

5. **Treat certification as a multiplier, not an endpoint**
   Continuous improvement turns a one off event into sustained value

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044178-1A

You can't wish for trust

In AI governance, hoping for the best isn't enough—you have to earn trust. ISO 42001 gave us a live fire test of our governance model, and the insights we gained apply to any organization navigating AI adoption.

Ask yourself: Could your governance withstand the gauntlet? And if not, what's your plan to get it there?

How KPMG can help

At KPMG, we help organizations design and build effective AI governance programs for the realities of today's increasingly autonomous systems. We support you in establishing practical policies, clear roles and accountabilities, and appropriate ethical guardrails—so governance becomes an enabler of safe, responsible adoption.

Whether you are starting your AI governance journey or maturing an existing program, we provide hands-on guidance across the AI lifecycle. This includes helping assess readiness against ISO 42001 requirements, supporting implementation of the controls needed to meet the standard, and assisting with certification efforts where applicable.

We also help you strengthen assurance and oversight through risk strategy, assessment, and readiness reviews aligned to evolving regulatory expectations. Our approach is designed to help you identify gaps early, address key security and privacy considerations, and put monitoring and remediation practices in place—so your AI ecosystem remains accountable, resilient, and secure as it scales.

KPMG

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. USCS044178-1A

Authors



Bryan McGowan
Principal, Advisory
Global & US Trusted AI Leader
KPMG US
E: bmcgowan@kpmg.com



Aisha Tahirkheli
Principal, BPG
Trusted AI & Product Risk
KPMG US
E: atahirkheli@KPMG.com



Kathi Craig
Sr Director, BPG
Trusted AI & Product Risk
KPMG US
E: kathiray@KPMG.com

Contributors

The authors thank Lisa Bigelow, Caleb Hargesheimer, Liz Kalooky, Sophia Pham, Gabriele Vaisnoraite, Olivia Weiss, and Kelvin Zlochevsky for their contributions to this paper.

References



Risk Modernization | The next chapter of operational risk: The AI rewrite



Cybersecurity considerations 2026



The trust advantage: Transforming at speed and scale

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: **kpmg.com**

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2026 KPMG LLP, a Delaware limited liability partnership, and its subsidiaries are part of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization. USCS044178-1A