KPMG

The Institute of Internal Auditors Singapore

The Agentic Opportunity: Governing AI for Trust, Integrity and Impact

A Practical Playbook for AI Organisational Transformation

KPMG

The Institute of Internal Auditors Singapore

Contents

- **Introduction** 03
- **Chapter 1** 04
  - 1.1 Setting the scene
  - 1.2 The three risk areas
- **Chapter 2** 06
  - 2.1 Human capital, talent & culture
  - 2.2 AI security
  - 2.3 Digital disruption & adaptive governance
- **Chapter 3** 10
  - 3.1 The action roadmap
  - 3.2 Connecting to the ecosystem
- **Conclusion** 12
- **Acknowledgements** 13

KPMG and The Institute of Internal Auditors Singapore would like to express our sincere appreciation to all contributors from various public sector, financial services, technology, and business services organisations who generously shared their time, expertise, and perspectives throughout the development of this publication.

Foreword

Artificial Intelligence (AI) is no longer a future prospect — it is reshaping how organisations operate, how employees work, and how value is created, right now. Singapore's national ambition as a Smart Nation is clear: to become a trusted global hub for AI, where innovation is matched by responsibility, and where the productivity gains of technology translate into better outcomes for workers and society alike. As AI becomes increasingly embedded in core decision-making and operating models, internal auditors — alongside business leaders, governance and risk professionals, and employees across the organisation — play a crucial role in shaping decision-making, managing risks, and driving value in AI adoption.

This playbook positions that organisations that govern AI responsibly — not simply adopt the most technology — will thrive, and that internal audit must step forward as a strategic partner in that effort.

Developed by KPMG and The Institute of Internal Auditors Singapore (IIAS), The Agentic Opportunity is a practical guide for business leaders navigating AI transformation. It draws on global risk survey data, case studies, and expert perspectives to address building an AI-ready workforce, securing AI systems, and governing AI in a landscape of constant change.

Singapore's national programmes — from workforce reskilling initiatives to world-leading AI governance frameworks — provide a strong foundation. This playbook shows how to build on it.

JONATHAN HO
Partner,
Head of Risk Consulting,
KPMG in Singapore

DAVID TOH
President,
The Institute of Internal Auditors Singapore

High-severity risks are outpacing internal audit coverage

The 2026 IIA Risk in Focus Singapore report highlights a nuanced picture: while cybersecurity shows strong alignment between perceived severity and audit coverage, the other two high-priority risks in human capital and AI-related digital disruption continue to outpace internal audit attention.

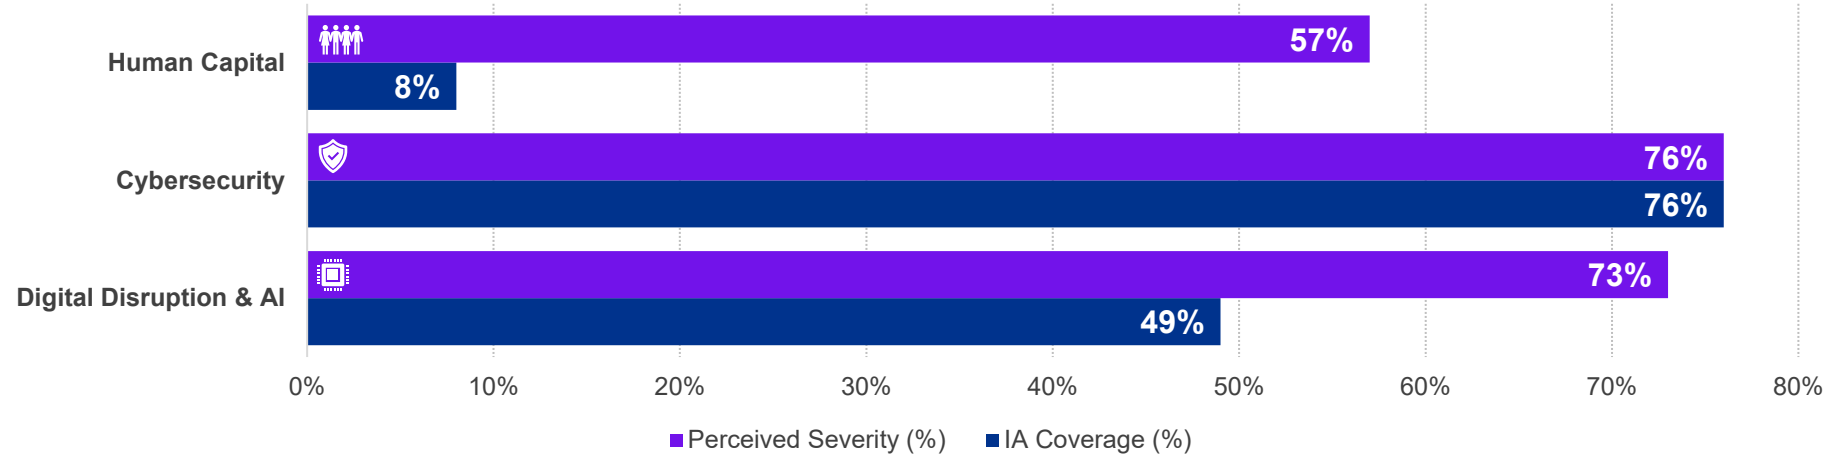

Organisations are aware of the risks. Yet, across all functions, including internal audit, they are not acting on them with sufficient urgency. This playbook is designed to close that gap.

> AI is raising new considerations for boards, particularly as they navigate a more complex risk landscape and heightened expectations around oversight. This creates an opportunity for internal audit functions to play a more active role in supporting directors — helping to build understanding of emerging risks and providing assurance on how these are being managed. Ultimately, this can strengthen the relevance of internal audit as organisations respond to the evolving impact of AI.

GERARD TOH
Partner, Risk Consulting
KPMG in Singapore

> Companies are still adapting to the role of AI, and there is a growing recognition that effective use of AI depends on keeping humans-in-the-loop to guide, interpret and validate outcomes. While AI can support specific tasks, it is not yet at a stage where it should operate independently without human oversight, which is crucial to ensure decisions are sound and risks appropriately managed.

ANDREW KOH
Partner, Technology Risk, Advisory,
KPMG in Singapore

AI transformation demands action across three interconnected areas

Area 1: Human Capital

The challenge

AI is reshaping how people work, shifting employees from execution to orchestration and increasing the need for human judgment and oversight. Yet many organisations remain culturally and structurally underprepared for this shift.

Key Solutions

- Build AI fluency at all levels of the organisation
- Avoid blanket bans, encourage safe adoption
- Rethink career pathways to progress in the age of AI
- Keep humans in the loop so executive decision remains with people

Area 2: Cybersecurity

The challenge

AI brings two overlapping vulnerabilities — AI risk, when models behave in ways that are unintended or hard to explain, and AI security, in which cyber-assailants seek to exploit AI systems for malicious purposes. Both must be governed in tandem.

Key Solutions

- Integrate AI risk and security into a single AI Trust framework
- Centralise all AI deployments
- Establish controlled AI deployment patterns (e.g. centralised, federated, or hybrid), aligned to risk appetite
- Evaluate model hosting strategies (public, private, hybrid, Retrieval-Augmented Generation-based) based on data sensitivity and use case risk
- Treat AI literacy as a first line of defence

Area 3: Digital Disruption

The challenge

Model drift — where AI performance degrades as datasets become outdated or incomplete — threatens the integrity of AI-enabled decisions. The human-in-the-loop principle is key to monitoring and identifying model drift and other risks.

Key Solutions

- Real-time monitoring over retrospective auditing
- Appoint Entitlement Gatekeepers for agentic AI
- Adopt official guidelines such as Singapore's 2026 Model AI Governance Framework for Agentic AI
- Set clear risk appetite at C-suite level

Across all three areas, internal audit must shift from post-implementation reviewer to proactive strategic adviser — involved from the design phase, not called in at the end.

Building AI fluency is not optional — it is the foundation of responsible governance

The Four-Stage Fluency Roadmap

STAGE 1 **Awareness**
Understand AI risks; learn to identify hallucinations; develop foundational data literacy.

STAGE 2 **Adoption**
Acquire operational AI skills; build personal AI assistants; write effective prompts.

STAGE 3 **Application**
Apply AI effectively in one's domain.

STAGE 4 **Amplification**
Scale AI use across teams and the enterprise; build shared AI ecosystems.

Building AI fluency is not optional — it is the foundation of responsible governance

Priority Actions

1 Build AI fluency
Equip staff to "trust but verify" through validation techniques such as source verification and cross-model checking, ensuring AI outputs are critically assessed rather than accepted at face value.

2 Avoid blanket bans on AI, encourage safe adoption
Banning AI drives shadow IT which can drive ungoverned workarounds. Set clear, well-communicated governance framework and usage guidelines.

3 Redesign career pathways for an AI future
Ensure professionals at all stages continuously build and refresh foundational problem-solving skills, even as AI automates routine tasks and reshapes traditional team structures.

4 Keep humans in the loop
Keep final authority for decisions with people and establish clear ownership for outcomes arising from AI-enabled processes.

Workforce Singapore Programmes

WSG supports job redesign, skills-first hiring, and career resilience pathways — including the **Career Conversion Programme** and the **SkillsFuture Workforce Development Grant** (Job Redesign+).

> Most organisations train their people to use AI — far fewer train them to question it. AI fluency must go beyond usage to include governance, data management, and the ethical challenges you will inevitably face. When things go wrong, internal audit should be the function that saw it coming, not the one called in to explain it.

TEA WEI LI
Partner, Risk Consulting,
KPMG in Singapore

> When it comes to decision-making, the human holds final authority – whether you are a board director or a business owner. AI can inform and accelerate, but it cannot be accountable. Leaders at every level need to ask whether they genuinely understand the AI-informed recommendations they are acting on, or are simply approving outputs they cannot interrogate.

GUILLAUME SACHET
Partner, Corporate Transformation,
KPMG in Singapore

AI security and AI risk must be governed together — not as separate silos

The Four Key Threats

1 Prompt Injection
Tricking AI into ignoring its own operating instructions to expose sensitive data (direct or indirect).

2 Data Poisoning
Corrupting training and reference data so AI learns to treat threats as safe (e.g. labelling a hacker's IP address as trusted).

3 Model Manipulation
Altering the way AI operates at the mathematical or structural level to bypass security controls (e.g. by fooling facial recognition).

4 Unauthorised Takeover of AI Agents
Using prompt injection to capture an AI agent and direct it to perform unauthorised actions — a growing risk as agentic AI spreads.

> AI does not only introduce incremental risks — it creates entirely new attack surfaces. An AI model can be manipulated through its inputs, corrupted through its training data, or hijacked to act against the organisation it was built to serve. Organisations of all sizes, including SMEs relying on third-party AI solutions, face comparable exposure as large enterprises. As AI becomes embedded in core business operations, a security breach will not be just another IT technical issue. It quickly escalates into concurrent operational, reputational, and regulatory risks and can result in severe consequences.

WENDY LIM
Partner, Cyber, Advisory,
KPMG in Singapore

The IA Role

Internal Auditors (IAs) must be strategic advisers — not just compliance auditors. They must prioritise AI audits as much as business audits. This means auditing the design and effectiveness of AI control frameworks, including guardrails, human-in-the-loop controls, as well as model validation and testing.

AI security and AI risk must be governed together — not as separate silos

> The worry is not a malicious user internally; it's an uninformed user that can put the organisation at great risk. So, you'd better train your users to know how to use AI responsibly. That's a soft guardrail, and you also need hard guardrails in your system.
>
> **SENIOR IA PROFESSIONAL**

> Because AI tools are so interesting, people are probably still using them outside the organisation. And probably putting company data out there. You can't just say 'don't use them'; make safe tools available for everybody.
>
> There is, of course, a need to track usage. If after six months usage is very low, we must review it to see whether we should kill some projects, move the funds and deploy other tools.
>
> **SENIOR IA PROFESSIONAL**

Hard & Soft Guardrails

Hard Guardrails

- Centralise all AI deployments through a unified hub with rigorous access controls.
- Use private LLMs to prevent corporate data leakage.
- Conduct adversarial testing for attacks such as prompt injection and model manipulation.
- Block external AI portals on company devices; make approved tools easy to use.
- Study SaaS vendor contracts to understand whether vendors use client inputs as training data.

Soft Guardrails

- Treat AI literacy as a first line of defence — employees who understand threats are harder to exploit.
- Promote human-in-the-loop as the most critical cybersecurity safeguard.
- Reject blame culture: treat AI failures as learning moments, not punishable offences.
- Build digital sandboxes for safe experimentation.

AI governance must be active, iterative, and human-led — not a static rulebook

The Threat of Model Drift

Model drift occurs when AI performance degrades because underlying datasets become outdated, incomplete, or contextually misaligned, or when the model itself becomes outdated. It is silent, gradual, and potentially catastrophic if undetected.

- **Machine learning:** Lower susceptibility to undetected drift. Outputs are structured and measurable, enabling quantitative monitoring. However, machine learning models are not immune — data drift, concept drift, and covariate shift remain significant risks requiring active controls.
- **Generative / agentic AI:** Higher overall risk and harder to govern. Drift can originate from multiple sources simultaneously — foundation model updates, prompt pattern shifts, retrieval degradation, and tool changes. Reasoning can be difficult to audit, increasing the risk of silent, compounding degradation.

Priority Response

Implement frequent monitoring — for example, by embedding auditing tools directly into the AI pipeline. Traditional retrospective sample testing is ill-suited for the speed and volatility of agentic workflows.

1 Entitlement Gatekeepers

Define who has authority to grant AI agents the ability to execute transactions and access sensitive databases. Treat AI agents as "privileged users" subject to strict identity and access management protocols. Agent-initiated actions must also be logged and reviewed.

2 Data Governance as Foundation

Use back-testing against historical data to check prediction accuracy. Consider testing AI with data known to contain exceptions, to ensure the tool identifies them. Other key aspects of data governance include data distribution monitoring and data quality checks.

3 Iterative, Dynamic Governance

Static policies become obsolete within 18 months. Governance frameworks must be continuously reviewed and synchronised with the rapid pace of AI innovation and regulatory frameworks.

> Trust in AI has to be built from the inside out — whether you are running a multinational or a ten-person business. Leaders at every level need to ask whether they have the right culture, the right people, and the right governance to use AI responsibly, and to remain accountable for its outcomes. These are not questions only large organisations need to answer — they are the foundation of responsible AI adoption at any scale.

GERRY CHNG
Partner, Cyber, Advisory,
KPMG in Singapore

AI governance must be active, iterative, and human-led — not a static rulebook

The IA Paradigm Shift

Traditional IA role	Strategic IA role in the AI age
▶ Post-implementation audit	Upstream involvement in AI design and governance
▶ Sample-based, retrospective testing	Real-time, AI-powered continuous monitoring
▶ Separate IT and business audits	Integrated AI Trust audit covering AI risks and AI security
▶ Reactive to incidents	Proactive risk identification and advisory
▶ Reports on technical uptime	Reports on AI decision integrity and commercial alignment

> Someone has to own every AI solution — accountable for ensuring the data is current, the outputs are correct, and the facts have been checked. This holds whether you are a large enterprise or an SME where one person manages everything. Internal audit adds real value here: not just checking outputs, but asking whether the right ownership and accountability structures exist in the first place.
>
> .

PAUL KENT
Partner, Corporate Transformation,
KPMG in Singapore

Horizon Watch

Governance frameworks must continuously evolve in tandem with emerging technologies. IAs will need to monitor and respond to technologies such as:

- **Quantum computing:** Will require post-quantum cryptography.
- **Artificial general intelligence:** General-purpose AI that trains itself.
- **Artificial super intelligence :** AI operating beyond human problem-solving capacity.

IA Advisory Role

Assess whether existing risk frameworks appropriately address emerging AI risks, including hallucination, accuracy, regulatory and accountability risks.

Ensure that the Audit Committee receives reporting not just on technical factors like uptime, but also on AI decision integrity and its alignment with commercial intent.

Internal audit must move upstream — from mostly assurance to proactive strategic counsel

> Regardless of the tool — whether it's a library, a search engine, or AI — whoever does the work owns the output. Accountability does not transfer to the technology; it stays with the person. Internal audit's role is to ensure that across organisations of every size, that accountability is never assumed — it is explicitly assigned.

EDMUND HENG
Partner, Technology Risk, Advisory,
KPMG in Singapore

> All my team members have a good understanding of how the data works in the organisation, so they know if the bot is hallucinating. They can challenge it. They can deep dive into the data fields and validate and sense check and fact check.

SENIOR IA PROFESSIONAL
Technology Sector

6 Priority Actions for IA Leaders

1. **Elevate to strategic adviser**
   Leverage end-to-end visibility to identify manual processes ripe for safe automation.
2. **Shift to upstream validation**
   Intervene in the design phase of AI solutions, before systems are fully implemented.
3. **Start with a "test of design"**
   Verify whether the organisation even has an AI governance framework before evaluating technical accuracy.
4. **Leverage AI to boost coverage**
   Transition from manual sample-testing to technology-enabled governance, with the prospect of using AI to audit AI.
5. **Bridge capability gaps**
   Second data scientists into the IA function to codify testing methodologies and transfer knowledge.
6. **Audit the prompt**
   Prompts must be audited in order to prevent unauthorised access to information.

Three stages to AI-ready governance — start where you are, build from there

1

Build the workforce, shift mindsets

- Shift employees from execution to orchestration of AI processes.
- Build a four-stage AI fluency roadmap: **Awareness** → **Adoption** → **Application** → **Amplification**.
- Restructure career pathways so junior staff still develop complex problem-solving skills.
- Create sandboxes for safe experimentation; encourage AI mentorships.
- Train staff to adopt a "trust but verify" mindset.
- Avoid blanket bans; set clear, unambiguous rules for AI use.

2

Prepare for the new AI reality

- Develop a unified AI Trust Framework (AI risk + AI security).
- Establish formal data governance with named role owners.
- Designate “Entitlement Gatekeepers” for agentic AI for human accountability.
- Implement real-time guardrails, automated alerts and kill switches.
- Route all AI deployments through a centralised hub.
- Adopt iterative, dynamic governance — review frameworks regularly.

3

Connect to the ecosystem

- Anchor governance policies to external standards (AI Verify, ISO 42001).
- Tap into Singapore's national support network: Workforce and Skills Singapore (WSSG), National AI Impact Programme (NAIIP), SMEs Go Digital, TechSkills Accelerator (TeSA).
- Mitigate third-party SaaS risk: audit vendor contracts; deploy private LLMs where possible.
- **For SMEs:** begin with low-cost, subscription-based AI tools to automate resource-intensive tasks. Adapt transformation strategies and governance expectations to their scale, capabilities, and constraints.

The IA imperative spans all three stages: IA must be involved from the outset and throughout the AI transformation journey.

> Traditional GRC frameworks were built for a world where risks are stable, controls are periodic, and accountability is clear — agentic AI breaks all three assumptions simultaneously. Whether you are revising existing frameworks or building governance habits from the ground up, internal audit must be at the table from the start. You can only catch small failures before they become crises if your oversight was designed to move at the speed of AI.

JONATHAN HO
Partner, Head of Risk Consulting,
KPMG in Singapore

Singapore's national AI infrastructure gives every organisation a head start

Government programmes

WSG's Career Conversion Programme (CCP)
Salary support to reskill mid-career workers across 30+ sectors, including audit and finance.

WSG's SkillsFuture Workforce Development Grant (Job Redesign+) (WDG (JR+))
Offers enhanced support for workforce transformation and job redesign through three key components: workforce consultancy, capability building initiatives and workforce tech solutions.

National AI Impact Programme (NAIIP)
Aims to make Singapore's workforce "AI Bilingual" by 2029 (from March 2026).

SMEs Go Digital (IMDA)
Funding for SMEs to adopt pre-approved automation solutions.

TechSkills Accelerator (TeSA)
Builds tech talent pipeline across sectors.

SkillsFuture AI Courses
Free subscriptions to premium AI tools for Singaporeans completing selected courses.

Governance frameworks & IA resources

Governance Frameworks

- **2026 Model AI Governance Framework for Agentic AI:** World's first specialised framework for autonomous AI systems.
- **AI Verify (IMDA):** Governance testing framework for AI performance validation.
- **ISO 42001:** International AI management system standard.
- **IMDA Open Innovation Platform:** Crowdsources digital solutions from a network of tech innovators.

IIA Resources

- **Global Practice Guide 'Developing a Risk-Based Internal Audit Plan':** Risk-based audit planning.
- **Cybersecurity Topical Requirement:** Guidance for auditing cybersecurity.
- **Global Technology Audit Guide:** Identify IT governance gaps.
- **Artificial Intelligence Knowledge Centre:** Curated resources on AI risks, governance, and auditing.
- **Artificial Intelligence Auditing Framework:** Practical framework for auditing AI systems and controls.
- **Vision 2035:** Global roadmap for the future of IA in a digital economy.

KPMG "You can with AI" Programme

End-to-end AI support from strategy to full-scale transformation, connected to KPMG's Trusted AI Centre of Excellence (from May 2026).

CONCLUSION

In the AI age, governance is not a constraint on innovation — it is the foundation of it

The organisations that will lead in the AI age are not necessarily those with the most advanced technology. They are those that **govern it most responsibly** — building trust with customers, regulators, and employees.

1

Treat AI governance as a strategic priority

Appoint accountable leaders, build living governance frameworks, and review them regularly.

2

Invest in your people as much as your technology

AI fluency, critical thinking, and human judgement are your most important AI risk controls.

3

Make internal audit your strategic partner

IA's end-to-end visibility and independence make it uniquely positioned to be the guardrail for responsible AI innovation.

> For some organisations, AI adoption is moving faster than governance frameworks can keep up. For smaller organisations, the bigger question isn't whether support exists; it's whether they are able to navigate the ecosystem and tap into the right resources. Having someone who understands the AI landscape playing an audit and advisory role is one of the most practical safeguards any organisation can put in place.

SARAH CHIN
Principal Advisor, Risk,
KPMG in Singapore

About KPMG and The Institute of Internal Auditors Singapore

About KPMG in Singapore

KPMG in Singapore is part of a global organization of independent professional services firms providing Audit, Tax and Advisory services. We operate in 138 countries and territories with more than 276,000 partners and employees working in member firms around the world. Each KPMG firm is a legally distinct and separate entity and describes itself as such. KPMG International Limited is a private English company limited by guarantee. KPMG International Limited and its related entities do not provide services to clients.

About The Institute of Internal Auditors Singapore

Internal auditing is an independent, objective assurance and advisory service designed to add value and improve an organisation's operations. It helps an organisation accomplish its objectives by bringing a systematic, disciplined approach to evaluate and improve the effectiveness of governance, risk management, and control processes.

As the only professional body dedicated to advancing the internal audit profession in Singapore, The Institute of Internal Auditors Singapore has been advocating for the interests of the profession for over 40 years. We remain committed to raising our profile, standing, and developing the knowledge, skills, and expertise of internal auditors. Representing more than 2,700 members, we enhance technical excellence through The IIA's new International Professional Practices Framework and Global Internal Audit Standards.

IIA Singapore is an affiliate of IIA Global. As affiliates, our members are concurrently members of IIA Global and have exclusive access to IIA Global's content, resources, and benefits. They belong to a global community of more than 270,000 professionals in over 170 countries and territories who share a common vision to advance their professional growth in internal auditing and add value to their organisations.

Contact us

Jonathan Ho
Partner, Head of Risk Consulting and Head of Infrastructure, Government & Healthcare
KPMG in Singapore
T: +65 6411 8336
jho1@kpmg.com.sg

kpmg.com.sg

The Institute of Internal Auditors Singapore
T: +65 6324 9029
E: secretariat@iia.org.sg

iia.org.sg