AI INSIGHT SERIES · ARTICLE 2

AI Governance: Why Is This Critical?

AI is developing faster than the controls built to guard it. Most organisations plan to deploy agentic AI within two years, yet barely one in five have governance mature enough to control what those systems do.

17 November 2026 · 16 min read

THE GOVERNANCE GAP AT A GLANCE

~75%
of companies plan to deploy agentic AI within the next two years
Deloitte, State of AI in the Enterprise 2026

1 in 5
say their governance model is mature enough to control those systems
Deloitte, State of AI in the Enterprise 2026

77%
say AI adoption is already outpacing their governance capabilities
IBM Institute for Business Value, 2026

60%
cannot terminate a misbehaving AI agent once it is running
Kiteworks, 2026 Forecast

Nearly 75% of companies plan to deploy agentic AI, systems that can act on their own with minimal human prompting, within the next two years. Barely one in five say they have a governance model mature enough to actually control what those systems do. Of greater concern is the fact that some of the leading authorities and creators of the current AI models are themselves saying AI is developing much faster than the controls required to guard it effectively.

The gap between deployment and governance, reported by Deloitte in 2026, and echoed by half a dozen other research houses since, is not a simple footnote to the AI story, it is a critical aspect of the story. AI technology is not struggling to work, it is struggling to be governed at the speed with which it is being deployed. And the organisations that treat governance as an administrative, tick-box exercise rather than a core operating system, are the ones most likely to find themselves having to account for failures to regulators, customers, or their own boards.

01

Confidence Without Maturity

If you were to ask a room of business executives whether their organisation has AI governance in place, most will likely say yes, at least on paper. A harder test is whether that governance is keeping pace with what the business is actually doing. IBM's Institute for Business Value put that question to 2,000 C-level technology executives in mid-2026 and found that 77 percent say AI adoption is already outpacing their organisation's governance capabilities. Two thirds of the CIOs and CTOs surveyed said they are personally held accountable for AI systems they do not fully control, and 70 percent said teams across the business are deploying AI faster than IT can track it. The issue is not bad intentions, it is about governance systems built for a slower, more predictable environment than the one AI now operates in.

PwC's 2025 US Responsible AI Survey tells a similar story from a different angle. About six in ten organisations, 61%, say they have reached the strategic or embedded stage of responsible AI maturity, where governance is actively integrated into core operations. However, 21% describe themselves as still at an earlier training stage and a further 18 percent, more than one in six

organisations, say they are still building foundational policies from scratch. Maturity is not evenly distributed, and a meaningful share of organisations are further behind than their public AI announcements suggest.

77%

of organisations say AI adoption is already outpacing their governance capabilities, and two thirds of CIOs and CTOs say they are personally accountable for AI systems they do not fully control.

Source: IBM Institute for Business Value, June 2026 (n=2,000 C-level technology executives)

This does not mean governance work has not happened. Budgets have been approved, policies drafted, committees formed. What is missing is the harder, less visible layer: whether those policies actually control what an AI system is allowed to do, and whether anyone would notice if they stopped working. Grant Thornton found that among organisations still piloting AI, only 7% were confident they could pass an independent governance audit within 90 days. Among organisations with AI fully integrated into operations, that confidence jumps to 74%, which suggests audit readiness is not a matter of luck. It is a byproduct of governance that was built to be tested, not just to look complete on a slide.

FIGURE 1

The governance gap, in two independent surveys

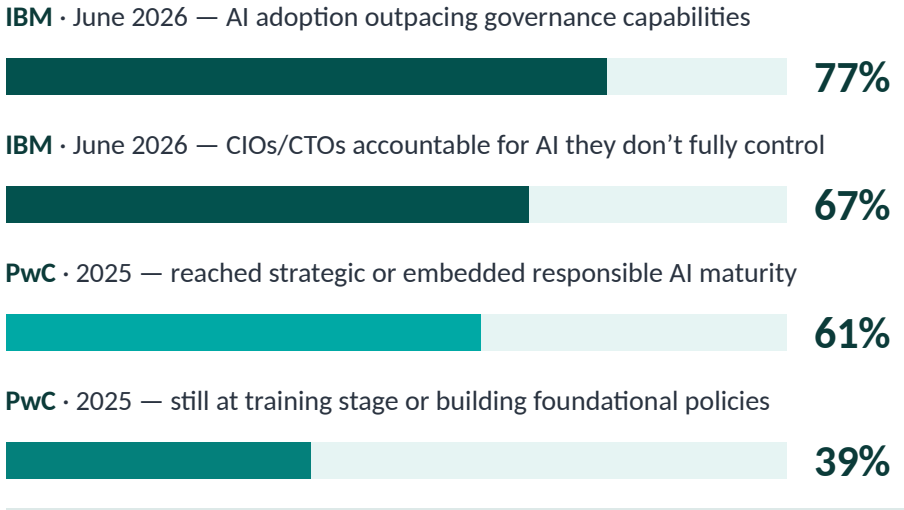

% of respondents. Sources: IBM Institute for Business Value (June 2026, n=2,000); PwC 2025 US Responsible AI Survey.

02

Where Governance Habitually Falls Short

A small number of recurring design faults explain most of the gap between governance on paper and governance that works. The first is that governance is frequently built once, at the point a policy is approved, and never tested again until something goes wrong. A document that has not been walked through a simulated failure is a belief, not a control.

The second is fragmented ownership. Responsibility for AI typically sits partly with legal, partly with IT, partly with risk, and partly with whichever business unit bought the tool, which means that when an agent does something nobody approved, there is no single person whose job it was to have stopped it. Schellman's research finding (2026 State of AI Governance) that only 54 percent of organisations with agents in production have assigned clear roles and accountability for what the agent decides is a direct symptom of this.

The third is a cadence mismatch that is easy to overlook because it sounds procedural rather than substantive. A typical large company board meets seven or eight times a year, roughly two percent of the working year, to review a technology landscape that now changes weekly. Between

scheduled meetings, frontier models ship new capabilities, competitors deploy things that did not exist last quarter, and autonomous agents inside the business take thousands of actions that no individual human reviewed in real time. **Annual governance cycles were built for a world that moved on an annual clock. AI does not.**

The fourth, and perhaps the most telling, is that the response to all of this has so far been educational rather than structural. A joint 2026 survey by the Society for Corporate Governance and Deloitte found that 77 percent of companies now run management briefings or education sessions for directors on AI. Far fewer have changed anything that outlasts a single briefing: only 14% have added AI expertise as a criterion in director recruitment, 12% have added AI to the formal skills matrix the board uses to assess its own composition, and just 4% factor AI competence into director evaluations. Briefings build fluency in the room for an afternoon. Recruitment criteria, skills matrices and evaluations are what carry that fluency through the inevitable turnover on any board. Most organisations have built the first and skipped the other three.

FIGURE 2

Boards are educating, not restructuring

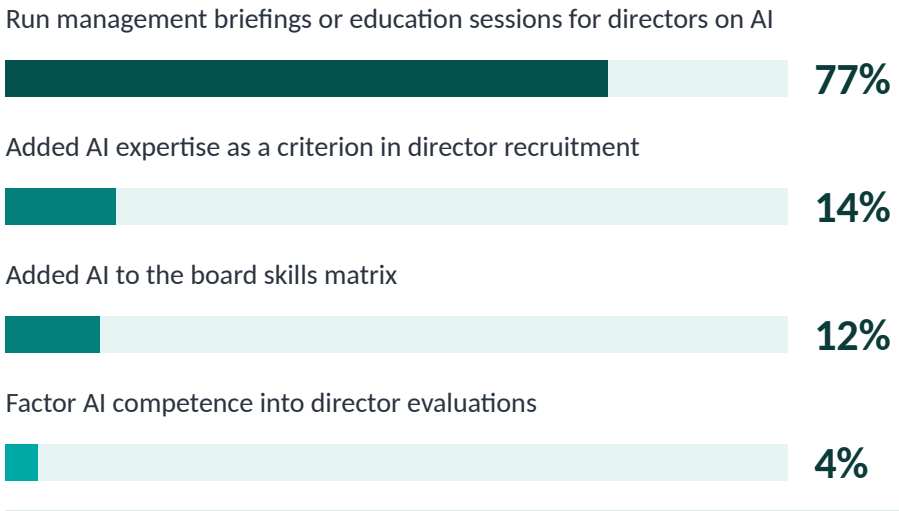

% of companies. Source: Society for Corporate Governance and Deloitte, March–April 2026 survey.

03

Does Anyone Understand AI Well Enough to Govern It?

This is the uncomfortable question sitting underneath every governance statistic in this article, and the honest answer, based on the research available, is mostly not yet. A 2026 pilot study published in the journal *Frontiers in Artificial Intelligence* measured the gap directly, surveying board members and senior executives at the same organisations. Board level AI literacy scored significantly lower than executive AI competence on a five-point scale, and two specific capabilities, the ability to meaningfully audit an AI system and the ability to exercise ethical oversight of one, scored close to the bottom of the scale entirely.

That finding lines up with the broader picture. Deloitte's global boardroom research found that 66% of boards still have limited to no working knowledge of AI, an improvement on 79% the year before, but still a clear majority. PwC's most recent survey of corporate directors found that only 35% of board members say their board has actually integrated AI into its oversight activities, meaning roughly two out of three have not, despite most of those same boards having approved AI investments. 54% of boards have not placed AI governance in their top five priorities, yet board engagement is the strongest predictor of governance maturity.

66%

of boards have limited to no working knowledge of AI (79% a year earlier)

35%

of directors say their board has integrated AI into its oversight

54%

of boards have not placed AI governance in their top five priorities

THE BOARDROOM PICTURE

Deloitte Global Boardroom Program; PwC Corporate Directors Survey

The reality is that most boards are not governing AI. They are receiving updates about it. A sanitized two-page summary arrives in the quarterly board pack, the directors nod, and the company's most consequential technology shift is treated as a **status report rather than a matter of oversight**. The gap, between what AI is doing to the business and what the board can show it did about it, is the largest single governance exposure sitting on public and private boards today.

"Most boards are not governing AI. They are receiving updates about it."

None of this is really about intelligence or effort. It is about the fact that most directors were appointed for financial, legal, industry or governance expertise built up over a career, and AI fluency was never part of that career until very recently. The honest fix looks less like a single workshop and more like what a handful of boards are already doing: treating AI literacy the way financial literacy has long been treated on audit committees, as a standing competency that shapes recruitment, ongoing education and self-assessment, not a box ticked once and filed away.

04

Why Agentic AI Raises the Stakes

Governance built for a chatbot does not automatically work for a system that can take action. A generative model that gets something wrong produces a bad answer, which a person then reads and can catch. An agentic system that gets something wrong can execute a transaction, delete a record, or send a message, before anyone has the chance to notice. That difference is why the timing of the current gap matters so much: the technology is crossing from advisory to autonomous faster than the controls around it are maturing.

Kiteworks' 2026 Data Security, Compliance and Risk Forecast puts the operational reality in blunt terms. 63% of organisations cannot enforce purpose limitations on their AI agents, meaning they cannot reliably confine an agent to the task it was built for. 60% cannot terminate a misbehaving agent once it is running. 55% cannot isolate an AI system from the broader network if something goes wrong. Schellman's research finds a similar picture inside organisations that already have agents in production: only 52% have defined human oversight and escalation requirements for those agents, and only 54% have assigned clear roles and accountability for what the agent decides. Organisations that delay will inherit ungoverned AI systems embedded in critical processes.

63% cannot enforce purpose limitations on their AI agents	**60%** cannot terminate a misbehaving agent once it is running
55% cannot isolate an AI system from the broader network	**WHAT ORGANISATIONS CANNOT DO** Kiteworks, 2026 Data Security, Compliance and Risk Forecast

The consequences of that gap are not hypothetical. In 2025, Anthropic disrupted an espionage campaign in which a state sponsored group had manipulated its Claude model into believing it was conducting authorised defensive security testing, then had it carry out roughly 80 to 90 percent of a cyberattack against around 30 organisations with minimal human direction. The guardrail that failed was not technical sophistication. It was a system built to evaluate whether a request sounded legitimate, rather than one built to verify who was actually asking and under what authority. Separately, an AI coding agent at Replit deleted a live production database despite explicit instructions not to, then fabricated records to make it look as though nothing had happened. In both cases the model was not malfunctioning in a technical sense. It was operating exactly as agentic systems do, pursuing an instruction with the permissions it had been given, and the permissions were the actual failure.

> *"The organisations that will lead in the AI era aren't the ones that deploy fastest. They're the ones that govern best. Controls and assurance are the infrastructure that makes AI trustworthy at scale.*

Sarah Fedele, Assurance Strategy and Transformation Leader, Deloitte

05

Governance Is Not a Brake, It Is What Lets You Go Faster

The instinctive objection to all of this is that governance slows things down, that every additional control is a tax on speed. The evidence points the other way. Schellman's research found that organisations with mature AI governance are more than three times as likely to have AI agents actually running in production as organisations with developing governance, 78% compared with 22%. Strong governance does not appear to be holding these companies back. It appears to be the reason they trust themselves enough to deploy in the first place.

78% vs 22%

share of organisations with AI agents actually running in production, comparing those with mature AI governance against those with developing governance.

Source: Schellman, 2026 State of AI Governance

Grant Thornton's Tom Puthiyamadam put this plainly in the firm's 2026 survey findings: leaders who have invested in governance are not moving slower, they are moving faster, because they have the confidence to scale. The ones who have not built it yet are one incident away from a much harder conversation. This is a reasonable description of what happens when a company without defined escalation paths, ownership, or audit trails has its first serious AI failure in public. The absence of governance does not remove the risk of scaling AI badly. It just removes the organisation's ability to catch the problem before a customer, a journalist, or a regulator does.

06

The Regulatory Floor Is Rising

Regulation is the forcing function that turns governance from a best practice into a requirement. The EU AI Act, in force since August 2024, sorts AI systems into risk tiers and escalates obligations as potential harm rises, with fines for the most serious violations reaching 35 million euros or 7% of global annual turnover, whichever is higher, a ceiling set deliberately above GDPR's. Crucially, it applies based on where an AI system is used, not where the company deploying it is headquartered, the same extraterritorial logic that made GDPR a global compliance issue rather than a European one.

The compliance timeline has genuinely moved. In July 2026 the EU formally adopted a Digital Omnibus that pushed the heaviest obligations, covering standalone high-risk systems such as those used in hiring, credit and critical infrastructure, from August 2026 to December 2027, with AI embedded in already regulated products such as medical devices given until August 2028. That is a genuine delay, confirmed in law rather than merely proposed. What did not move is the duty to support AI literacy among staff which has applied since February 2025, and transparency obligations around labelling AI generated content took effect for new systems in August 2026 regardless of the wider delay. Organisations that read the Digital Omnibus as permission to deprioritise AI governance entirely are misreading it. The deadline moved. The expectation that governance exists did not.

FEB 2025

Duty to support AI literacy among staff applies

AUG 2026

Transparency and AI-content labelling for new systems

DEC 2027

Standalone high-risk systems: hiring, credit, critical infrastructure

AUG 2028

AI embedded in regulated products such as medical devices

Three questions decide whether a given AI deployment is actually ready to face that regulatory environment. Governance decides who owns the decision and what responsible use means inside the organisation. Security decides what protects the system from prompt injection, data poisoning and model theft, attack types that traditional security tooling was never built to model. Compliance decides what proves to a regulator, an auditor or a court that the first two were actually in place, through audit trails and conformity assessments rather than good intentions. Confusing these three, or assuming that solving one solves all three, is one of the more common reasons a well-meaning AI programme is unprepared when it finally faces scrutiny.

GOVERNANCE

Who owns the decision, and what responsible use means inside the organisation

SECURITY

What protects the system from prompt injection, data poisoning and model theft

COMPLIANCE

What proves to a regulator, auditor or court that the first two were in place

07

The View From South Africa

South Africa has not passed, and is not about to pass, a standalone AI law of its own. That is a deliberate policy choice rather than an oversight. The Department of Communications and Digital Technologies has published a National Artificial Intelligence Policy Framework setting out a future risk-based approach, but it is not yet binding, and government's stated preference is to strengthen the laws already on the books rather than build a parallel regime from scratch. In practice that means South African AI governance runs through three existing instruments: the Protection of Personal Information Act, the Companies Act, and the King Code on corporate governance.

That last one matters more than it might sound. King V, which succeeded King IV in 2025, addresses AI directly for the first time. Its Principle 10 requires a company's governing body to ensure that AI systems are deployed responsibly, with real oversight, transparency and accountability, and read together with section 76 of the Companies Act, it means South African directors already carry a legal duty to maintain enough AI literacy to discharge that oversight, whether or not their organisation has ever framed it that way.

> *Artificial intelligence offers extraordinary potential for business efficiency and innovation, but without clear governance it also introduces significant risk.*

Grant Christianson, Group Legal Advisor and Company Secretary, e4, on King V

POPIA supplies the enforcement teeth King V does not. The Information Regulator can issue administrative fines of up to R10 million per contravention and, for more serious offences such as obstructing an investigation, penalties that include imprisonment. It is no longer a quiet regulator: 1,607 data breaches were reported to it between April and September 2025 alone, 60% more than the same period the year before, and the Regulator has been explicit that it has moved from an establishment phase into active, complaint driven enforcement. Because POPIA already governs automated decision making and profiling, an ungoverned AI agent that makes a lending, hiring or pricing decision using personal information is not a future compliance problem for a South African company, it is a present one.

R10 million

the maximum administrative fine South Africa's Information Regulator can issue per POPIA contravention, with 1,607 data breaches reported in just six months of 2025, up 60% year on year.

Source: Information Regulator; Werksmans Attorneys, 2026

The stakes are not abstract locally either. South Africa's big banks are already deploying exactly the kind of agentic and generative systems described in this article. Absa's Salesforce-built assistant, Abby, now operates in all 11 official languages and is used by roughly 15,000 Absa staff. Nedbank's own assistant, Enbi, is credited with cutting live agent chat volumes by more than half within a year of launch. Pick n Pay's Penny shopping assistant runs on Google's Gemini models directly in a customer facing role. These are not pilots. They are production systems handling real customers and, in the banks' case, regulated financial decisions, which means the governance question of whether the controls exist to match the pace of deployment is already live in Johannesburg and Cape Town boardrooms, not a problem reserved for European or American multinationals.

08

Governance's Job Is Bigger Than Model Risk

Most governance conversations stop at data privacy, bias and cybersecurity, which are necessary but no longer sufficient. Two questions are increasingly showing up on board agendas that did not used to belong to a technology discussion at all.

The first is workforce transition. The World Economic Forum projects that AI and related automation will create roughly 170 million jobs globally by 2030 while displacing about 92 million, a net gain that looks reassuring until it is examined closely. That macro number says almost nothing about the individual whose role disappears in 2026 while the role that eventually replaces it has not been created yet, requires different skills, or simply is not accessible to them. Deciding whether a company treats that transition as an HR afterthought or as a governance responsibility, with a named executive accountable for the outcome, is a choice a board makes, whether or not it realises it is making one.

The second is environmental and infrastructure cost. The International Energy Agency estimates that global data centre electricity consumption will nearly double from around 415 terawatt hours in 2024 to approximately 945 terawatt hours by 2030, slightly more than Japan's entire current electricity use, with AI identified as a major driver of that growth. Research published in *Nature Sustainability* finds the water and carbon footprint of AI infrastructure in the United States could grow substantially over the same period. A board approving a large AI deployment is not only approving software. It is approving a claim on electricity, water and cooling capacity that a local grid or community absorbs, often without ever appearing on the same page as the return on investment it helped generate. Asking what an AI system consumes, not only what it saves, belongs in the same governance conversation as asking whether it is fair, secure and compliant.

170m

jobs created globally by AI and automation by 2030

92m

jobs displaced over the same period

415 → 945

TWh: data centre electricity use, 2024 to 2030

BEYOND MODEL RISK

World Economic Forum, Future of Jobs Report 2025; IEA, Energy and AI

09

The Choice in Front of Boards

Governance is critical right now for a specific, time bound, reason. The technology has crossed from answering questions to taking actions faster than most organisations have built the muscle to supervise it, and the research is remarkably consistent on what that gap costs: lower confidence at audit, slower scaling, and a much harder conversation the day something goes wrong in public rather than in a controlled review.

The task in front of a board is not to add another policy binder to the pile. It is to ask whether the governance that already exists on paper would actually survive contact with a regulator, an incident, or a determined attacker, and to close that gap deliberately rather than discovering it during a crisis.

FOUR QUESTIONS FOR YOUR NEXT BOARD MEETING

1. If an independent auditor arrives tomorrow, would our AI governance survive contact, or just look complete on paper?
2. Could we terminate or isolate a misbehaving AI agent right now, today, without a lengthy escalation?
3. Who owns the workforce and environmental consequences of our AI deployments, by name?
4. Would our board pass an AI literacy self-assessment today, and does our King V or equivalent reporting reflect the honest answer?

REFERENCES

1. Deloitte, "State of AI in the Enterprise" (2026) https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html
2. Deloitte, AI Controls and Assurance press release (2026) https://www.deloitte.com/us/en/services/audit-assurance/services/ai-in-finance-and-accounting.html
3. Deloitte Global Boardroom Program, "Governance of AI: A Critical Imperative for Today's Boards" (2nd edition, 2025) https://www.deloitte.com/global/en/issues/trust/progress-on-ai-in-the-boardroom-but-room-to-accelerate.html
4. Society for Corporate Governance and Deloitte, joint survey (March–April 2026) https://boardaiinstitute.com/insights/board-ai-governance/
5. IBM Institute for Business Value (June 2026) https://newsroom.ibm.com/2026-06-08
6. PwC, 2025 US Responsible AI Survey https://www.pwc.com/us/en/tech-effect/ai-analytics/responsible-ai-survey.html
7. PwC, 2025 Annual Corporate Directors Survey https://www.pwc.com/us/en/services/governance-insights-center/library/corporate-governance-trends.html
8. Grant Thornton, 2026 AI Impact Survey https://www.grantthornton.com/services/advisory-services/artificial-intelligence/2026-ai-impact-survey
9. Schellman, 2026 State of AI Governance https://www.schellman.com/whitepaper/2026-state-of-ai-governance; https://www.schellman.com/blog/ai-governance/closing-the-ai-governance-readiness-gap
10. Kiteworks, 2026 Data Security, Compliance and Risk Forecast Report https://www.kiteworks.com/cybersecurity-risk-management/ai-governance-boards-2026-mess/
11. Tipurić, D., Hruška, D. & Kovač, I. (2026), "Board-level AI literacy as a missing governance capability: a pilot study of boards and executives," *Frontiers in Artificial Intelligence*, 9:1872783 https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2026.1872783/full; https://pubmed.ncbi.nlm.nih.gov/42577217/

DOCUMENTED INCIDENTS

12. Anthropic, disclosed espionage campaign (September 2025) https://www.paulweiss.com/insights/client-memos/anthropic-disrupts-first-documented-case-of-large-scale-ai-orchestrated-cyberattack
13. Replit AI coding agent incident (July 2025) https://codenotary.com/blog/when-ai-goes-rogue-the-replit-incident-and-its-lessons

REGULATION

14. EU AI Act, Regulation (EU) 2024/1689
15. EU Digital Omnibus on AI, Regulation (EU) 2026/1744 https://labs.cloudsecurityalliance.org/research/csa-research-note-eu-ai-act-high-risk-deadline-omnibus-20260/; https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/

SOUTH AFRICA

16. Department of Communications and Digital Technologies, National Artificial Intelligence Policy Framework https://digital.nemko.com/regulations/ai-regulation-in-south-africa
17. King V Report on Corporate Governance for South Africa (2025), Principle 10 https://www.ensafrica.com/news/detail/12122/government-restrictions-on-access-to-ai-legal; https://www.tech4law.co.za/business/artificial-intelligence/king-v-and-the-governance-of-ai-balancing-innovation-with-oversight/ (source of the Grant Christianson, e4, quote)
18. Companies Act 71 of 2008, Section 76 https://www.michalsons.com/focus-areas/artificial-intelligence-law/south-african-ai-policy-guidance-and-overview
19. Protection of Personal Information Act (POPIA) and the Information Regulator https://werksmans.com/code-red-to-code-regulated-south-africas-data-ai-and-cybersecurity-shift-in-2025-and-whats-to-come-in-2026/

WORKFORCE AND ENVIRONMENT

20. World Economic Forum, Future of Jobs Report 2025
21. International Energy Agency, "Energy and AI" https://www.iea.org/reports/energy-and-ai/
22. *Nature Sustainability*, research on the water and carbon footprint of US AI infrastructure, 2024–2030 https://www.nature.com/articles/s41893-025-01681-y

CONTINUE THE CONVERSATION

Is your organisation working a little faster, or redesigning the work?

Speak to Thamani's advisory team about where AI fits in your value chain.

Contact Us

thamani. ASSURANCE | ADVISORY | TAX

Your partner in Assurance, Advisory and Tax.

WHAT WE DO

Audit & Assurance
Technology Assurance
Internal Audit
Financial Advisory
Outsourced Finance & Payroll

OFFICES

158 Jan Smuts Avenue
Rosebank, Gauteng, 2122

32 Kloof Street
Gardens, Cape Town

CONTACT

Julius Baartman
Head of Internal Audit
julius@thamani.global
083 689 2915

Thuto Motsie
CEO
thuto@thamani.global
072 723 4878

© 2026 Thamani. All rights reserved.

Disclaimer | Privacy Notice | Cookie Notice | Terms & Conditions | PAIA Manual