EXECUTIVE GUIDE

The Executive's Guide to AI Governance

A practical framework for South African boards.

By Zoë François

Founder, The AI Fluency Company

First edition · 2026

CONTENTS

What's in this guide.

FOREWORD

Why this guide exists.

AI is not a future board issue. It is a present one.

In the past year, I have had many conversations with senior leaders — chief executives, finance directors, operations heads — and listened to essentially the same conversation. Their organisations are using AI. Sometimes deliberately, often by accident. Tools they bought for one purpose now have AI features they didn't ask for. Teams have started experimenting with various AI platforms and many elements of shadow AI systems are showing up. Vendors are pitching AI-powered everything. And nobody around the boardroom table has a clear view of what is happening, where the risks are, or how to govern any of it.

The information available to leaders is not always helpful. It is too technical, or too vague, or too vendor-driven, or too fearful. Compliance teams talk about ISO standards and regulators in language that requires a lawyer to translate. Technology teams talk about models and pipelines in language that requires an engineer to translate. Consultants pitch transformation programmes that assume a maturity most organisations don't yet have.

This guide is the resource I wanted to be able to hand to those leaders — a single document, written in their language, that gives them a working understanding of AI governance without asking them to become specialists. It is grounded in ISO/IEC 42001:2023 — the international standard for AI management systems, published in December 2023 and now becoming the framework most regulators reference — because that standard is the closest thing the world has to a settled answer about what good AI governance looks like. But it is written for the senior generalist, not the certification auditor.

> "The organisations that adopt AI well are not those with the most sophisticated technology. They are those with the clearest governance — the discipline to know what they are using, why, and what they would do if it went wrong."

The South African context matters here. The National AI Policy Framework was published in August 2024. A Draft National AI Policy was gazetted on 10 April 2026 (Government Gazette Notice 3880 of 2026) for public comment — and then withdrawn just 16 days later, on 26 April 2026, after journalists discovered that at least six of the document's 67 academic citations were entirely fabricated. The journals were real. The articles were not. Minister Solly Malatsi confirmed the withdrawal, stating that the inclusion of fictitious sources was not a technical glitch but a failure of human oversight over AI-generated content. The irony is not subtle: South Africa's AI governance policy was undone by the very governance failure it was designed to

prevent. Full implementation of a revised policy is now targeted for 2027 to 2028, though that timeline may slip. The standards being referenced in those documents map closely to ISO/IEC 42001:2023. South African organisations that build their AI governance around the standard now will be ready for the regulatory landscape that is coming — not surprised by it.

This guide will not make you an AI specialist. That is not its purpose. It will give you a frame for the conversations you need to be having — with your stakeholders - board, your executive team, your employees, your auditors, your customers — about the AI that is already in your business. And it will help you decide what to do next.

Zoë François

Cape Town, 2026

THE HISTORY OF AI

100 YEARS OF INNOVATION. ENDLESS POTENTIAL.

From early ideas to intelligent systems that power our world – AI has evolved through bold breakthroughs, setbacks, and unimaginable progress. The next 25 years will redefine what's possible.

1950s — THE BIRTH OF AI

1950
Alan Turing introduces the Turing Test in "Computing Machinery and Intelligence"

1956
The term "Artificial Intelligence" is coined at the Dartmouth Conference

1960s — EARLY EXPLORATION

1966
ELIZA, one of the first chatbots, is created

1969
Shakey the Robot performs basic problem-solving and navigation

1970s–1980s — AI WINTERS

1970s–1980
Limited funding and high expectations lead to the first AI winter

1987–1993
The second AI winter brings further slowdown in progress

1990s — FOUNDATIONS FOR GROWTH

1997
IBM Deep Blue defeats chess champion Garry Kasparov

1990s
Advances in computing power and data lay the groundwork for AI growth

2000s — MACHINE LEARNING TAKES SHAPE

2001
IBM Watson wins Jeopardy!

2006
Deep learning gains traction with Geoffrey Hinton's breakthrough research

2010s — DEEP LEARNING REVOLUTION

2012
AlexNet wins ImageNet competition, marking a turning point for deep learning

2016
AlphaGo defeats world Go champion Lee Sedol

2020s — AI GOES MAINSTREAM

2020
Large language models (LLMs) transform the way we work, learn, and create

2023
Generative AI goes mainstream with tools like ChatGPT, DALL·E, and beyond

2030s–2050 — THE NEXT ERA OF INTELLIGENCE

The next 25 years will bring:

- Autonomous systems at scale
- Human-AI collaboration at unprecedented levels
- Scientific breakthroughs accelerated by AI
- Ethical, responsible, and inclusive AI by design
- AI as a partner in solving humanity's greatest challenges

2050
100 YEARS OF AI
A future limited only by our imagination.

KEY THEMES THROUGH THE DECADES

IDEAS & IMAGINATION
Bold ideas spark the beginning

EXPERIMENTATION & DISCOVERY
Early research tests the limits

CHALLENGES & REALITY CHECKS
Setbacks shape a stronger foundation

TECHNOLOGY BREAKTHROUGHS
Computing power and data unlock new possibilities

WIDESPREAD APPLICATIONS
AI moves from labs to real-world impact

TRANSFORMATION AT SCALE
AI becomes part of everyday life and work

A FUTURE WE BUILD TOGETHER
The future is ours to shape responsibly

THE AI FLUENCY CO | EVALUATE. ELEVATE. EMPOWER.

CHAPTER 1

The governance imperative.

Three forces are converging — and they meet in your boardroom.

The question of whether AI matters at the board level is no longer interesting. It is imperative. The interesting questions are: how much, how soon, and what should we do about it? Three forces, working in parallel, are turning AI from a technology question into a governance one.

Force one: AI is already in your business.

Most organisations underestimate by an order of magnitude how much AI is already operating inside their walls. This is not because they are inattentive. It is because AI has become the default. It is built into the productivity software your office uses. It powers features in your customer relationship system, your finance platform, your video conferencing tool. Your customer service team is using it. Your sales team is using it. Your communications team is using it. Its not new AI has been around for years!

In a typical mid-market business today, you can name a dozen AI applications without doing any new work — they are already running. The board's job is not to decide whether to adopt AI. That decision has been made, by hundreds of small purchasing decisions over the past three years. The board's job is to decide whether to govern what is already there and manage what comes next.

> "The first audit any AI governance programme should do is an audit of what already exists. Most organisations are surprised by what they find."

Force two: stakeholders are starting to ask.

For the past two years, AI has been a curiosity — something boards heard about, occasionally read about, sometimes talked about in passing. That phase is ending. Boards, customers, regulators, and employees are now beginning to ask the same question, in different language: where is AI in our business and how is it governed?

Investors ask in due diligence. Auditors ask in their annual reviews. Major customers ask before they renew. Procurement teams ask before they sign. Insurance underwriters are starting to ask. And employees are increasingly asking about how AI affects their work, their performance reviews, their data and future.

Most organisations cannot yet answer these questions cleanly. That gap — between the question being asked and the organisation's ability to answer it — is the core risk that AI governance addresses. The cost of being unable to answer is not abstract. It shows up in lost contracts, delayed deals, escalated audit findings, and customer concern.

Force three: the regulatory clock is running.

South African policy on AI moved decisively in 2024. The National AI Policy Framework was published in August of that year by the Department of Communications and Digital Technologies, signalling government's intent to develop a comprehensive regulatory approach. On 10 April 2026, the Draft National AI Policy was published in the Government Gazette (Notice 3880 of 2026) for public comment — ambitious in scope, proposing five new oversight bodies including a National AI Commission. Sixteen days later, it was withdrawn. Journalists checking the bibliography discovered that at least six of the document's 67 academic citations were entirely fabricated. The journals cited were real. The articles were not. Authors were credited with foundational research on AI governance they had never written. Minister Solly Malatsi withdrew the draft on 26 April 2026, confirming that the document's integrity had been compromised by AI-generated content that was never verified by human oversight. "This unacceptable lapse proves why vigilant human oversight over the use of artificial intelligence is critical," he stated. The irony is not subtle: South Africa's AI governance policy was undone by the very failure it was designed to prevent. A revised policy has no confirmed timeline. Full implementation is now targeted for 2027 to 2028, though this may slip further.

The international picture is moving faster. The European Union's AI Act came into force in August 2024, with high-risk obligations applying from August 2026. ISO/IEC 42001:2023 — the international standard for AI management systems — was published in December 2023 and is being referenced by regulators globally. The UK has issued AI assurance guidance. The OECD has published AI principles that most member countries are aligning to. The standards being referenced across these jurisdictions are converging — and they look a lot like ISO/IEC 42001:2023.

For South African organisations, this matters in two ways. First, your local regulatory horizon is roughly two years out, and you cannot build an AI governance programme overnight. Second, if you trade internationally — sell into European markets, audit to global standards, work with multinational customers — you are already inside the regulatory perimeter of those jurisdictions. The clock is running on both fronts.

WHAT THIS MEANS FOR THE BOARD

AI governance is not optional and it is not delayable. The forces driving it are already in motion and accelerating.

The cost of acting now is modest. The cost of acting after a regulatory event, a customer escalation, or an AI incident is large.

The good news: the framework for what to do is settled. ISO/IEC 42001 gives you a clear, internationally recognised path. The work is not to invent governance — it is to apply governance you already know how to apply, to a domain that is new.

CHAPTER 2

ISO/IEC 42001, in plain English.

A management system standard for AI — written for technical specialists, but it doesn't have to be read that way.

ISO/IEC 42001:2023 — to give it its full name — is the world's first management system standard for artificial intelligence. It was published by the International Organization for Standardization in December 2023, after several years of development by an international working group. It is now being adopted as the foundation for AI governance frameworks in jurisdictions across the world, including South Africa.

The full standard runs to several dozen pages of formal language. It is written, like all ISO standards, in a way that is precise but not always accessible. This guide will not reproduce that language. Instead, it will give you what a senior leader needs to know about the standard — what it is, what it requires, and how it shapes the rest of this guide.

What it is.

ISO/IEC 42001 is a management system standard. That phrase has a specific meaning in the ISO world. It does not specify what AI you should use, or how AI models should be built, or what is and isn't allowed. It specifies the management system you should have in place to govern AI — the policies, processes, controls, and accountabilities that ensure AI is being used responsibly across your organisation.

If you have ever been involved in ISO 9001 (quality management), ISO 27001 (information security), or ISO 14001 (environmental management), you already know the shape of this. ISO/IEC 42001:2023 is the same shape, applied to AI. It specifies what your AI Management System (AIMS) must contain, how it must operate, and how it can be audited and certified.

The fact that it is certifiable matters. Certification is what turns a good intention into something you can demonstrate to a stakeholder. An organisation that says "we take AI governance seriously" is making a claim. An organisation that holds an ISO/IEC 42001:2023 certification has had that claim tested by an accredited third party.

What it requires — the seven core areas.

ISO/IEC 42001 organises its requirements into seven core areas. You don't need to memorise these. But it helps to know they exist, because everything else in this guide — and everything any auditor or regulator will eventually ask you about — maps onto them.

AREA	WHAT IT COVERS
Context of the organisation	Understanding your AI environment, the stakeholders affected by your AI use, and the scope of your AI Management System.
Leadership	Top management commitment to AI governance, with documented policies and clear roles.
Planning	Identifying AI-related risks and opportunities, and setting objectives to address them.
Support	Ensuring you have the people, skills, awareness, communication, and documentation to make AI governance work.
Operation	Implementing controls across the AI lifecycle — from concept through decommissioning.
Performance evaluation	Monitoring, measuring, auditing, and reviewing whether your AI Management System is working.
Improvement	Acting on what you learn, correcting issues, and continually improving the system.

> "ISO/IEC 42001:2023 doesn't teach you how to build AI. It ensures you don't lose control of it."

Why it matters now.

Three reasons. First, because boards and stakeholders are asking the questions ISO/IEC 42001:2023 is designed to answer. Adopting the framework means you can answer them — and answer them in language regulators and auditors recognise.

Second, because South African regulation is converging on these standards. The 2024 National AI Policy Framework and the 2026 draft policy reference the same governance dimensions. Building your AIMS now puts you ahead of the regulatory curve, not behind it.

Third, because the discipline of building an AI Management System surfaces the things you didn't know you didn't know. Most organisations begin the work assuming their AI footprint is small and their risks are

manageable. By the time they finish a serious assessment, they have a more accurate picture — and the foundation for managing it.

FOR SOUTH AFRICAN BOARDS: KING IV, KING V AND ISO/IEC 42001:2023

South African boards are already subject to governance obligations that map directly onto ISO/IEC 42001:2023. King IV's Principle 1 (ethical and effective leadership), Principle 11 (governance of risk), and Principle 15 (assurance) are all directly applicable to AI governance. Boards that apply these principles to AI — defining ethical boundaries, integrating AI into enterprise risk management, and commissioning independent assurance of AI systems — are already doing what ISO/IEC 42001:2023 requires.

The IoDSA released King V in November 2025 — the most significant governance reform in nearly a decade. King V goes further than King IV on technology governance, explicitly requiring boards to govern data, cyber, and AI as strategic risks, with public disclosure obligations. Organisations must now not only govern AI but demonstrate that governance through a published Disclosure Framework. ISO/IEC 42001:2023 certification is the most direct path to satisfying King V's AI disclosure requirements with internationally recognised evidence. Boards that align to the standard now will be ahead of both the regulatory curve and their governance obligations. Those that wait will face both simultaneously.

HOW TO READ THIS GUIDE FROM HERE

The next chapter takes the seven core areas of ISO/IEC 42001:2023 and translates them into ten practical pillars that frame how a senior leader should think about AI governance. Each pillar is a chapter in itself.

You don't need to read the pillars in order. They map onto each other but also stand alone. Use them as a reference for whichever conversation you are about to have with your board, your customers, or your team.

If you read nothing else, read the box at the end of each pillar titled 'Three questions to ask your organisation'. Those questions are the test of whether your governance is real or theoretical.

THE AI FLUENCY CO
EVALUATE. EDUCATE. EMPOWER.

The 10 Pillar Framework

AI POLICY OVERLAYS THE ENTIRE SYSTEM – CONNECTING STRATEGY TO TRUST

AI Policy is the foundation that sets direction, defines boundaries and drives consistent, responsible and accountable AI across the full lifecycle.

BUILD TRUST.
MANAGE RISK.
DELIVER VALUE.

AI POLICY – THE CONNECTIVE LAYER

Defines intent. Sets rules. Drives controls. Enables trust.

OUR COMMITMENT
We use AI to create value, responsibly and ethically.

OUR PRINCIPLES
Fairness, transparency, accountability, privacy, safety and inclusivity.

OUR RULES
What we allow, what we require, what we never do.

OUR ACCOUNTABILITY
Clear roles, governance and oversight at every level.

1 STRATEGIC INTENT
(Why AI Exists)

- AI vision aligned to business strategy
- Defined objectives & success criteria
- Use case prioritisation
- Risk appetite & ethical stance

AI with purpose. Value with direction.

2 GOVERNANCE & LEADERSHIP
(Who is Accountable)

- Leadership commitment
- Defined roles & responsibilities
- Oversight structures
- Policies (AI, ethics, data, security)

Accountability drives responsible AI.

3 RISK & IMPACT MANAGEMENT
(What Could Go Wrong)

- Risk identification (bias, privacy, safety, misuse)
- Impact assessments (legal, ethical, societal)
- Risk evaluation & prioritisation
- Mitigation & controls

Understand risk. Reduce impact.

4 AI LIFECYCLE MANAGEMENT
(How AI is Controlled)

- Concept & Use Case Definition
- Design & Development
- Validation & Testing
- Deployment
- Operation & Monitoring
- Maintenance & Improvement
- Decommissioning

Control every stage. Add value always.

5 DATA & TECHNOLOGY GOVERNANCE
(What it Runs On)

- Data quality & integrity
- Data lineage & provenance
- Model transparency & explainability
- Infrastructure & security controls

Quality data. Trusted systems.

6 HUMAN OVERSIGHT
(Who Stays in Control)

- Human-in-the-loop design
- Escalation & intervention mechanisms
- Decision accountability
- User training & awareness

AI supports. Humans decide.

7 PERFORMANCE EVALUATION
(Is it Working?)

- Monitoring KPIs & outcomes
- Bias & fairness tracking
- Drift detection
- Audit & review processes
- Measure ROI

Measure performance. Build confidence.

8 CONTINUAL IMPROVEMENT
(Does it Get Better?)

- Incident learning
- Model retraining
- Process optimisation
- Corrective & preventative actions

Learn continuously. Improve always.

9 COMPLIANCE & ASSURANCE
(Can You Prove It?)

- Documentation & records
- Internal audits
- Regulatory alignment
- Certification readiness

Evidence trust. Demonstrate compliance.

10 STAKEHOLDER TRUST & TRANSPARENCY
(Why People Believe You)

- Transparency of AI use
- Ethical disclosures
- User communication
- Trust frameworks

Transparency builds trust. Trust drives adoption.

AI POLICY IN ACTION ACROSS ALL PILLARS
Sets the intent & boundaries
Defines roles & accountability
Sets risk rules & thresholds
Defines lifecycle controls & gates
Sets data & tech standards
Defines human oversight rules
Defines what we measure
Requires improvement
Provides the evidence base
Builds transparency & trust

THE OUTCOME →

PROTECT PEOPLE & DATA
Manage risk and safeguard what matters.

ENSURE ETHICS & FAIRNESS
Build AI that is fair, inclusive and responsible.

DRIVE VALUE & PERFORMANCE
Deliver measurable outcomes aligned to business goals.

EARN TRUST & CONFIDENCE
Strengthen trust with users, regulators and stakeholders.

ENABLE INNOVATION WITH CONTROL
Innovate faster with governance that scales.

AI Policy is not just documentation.
It is how leadership exerts control over intelligent systems.

THE AI FLUENCY WAY

EVALUATE
Make the right AI decisions

ELEVATE
Build AI the right way

EMPOWER
Use AI to lead, not just keep up

"ISO 42001 doesn't teach you how to build AI. It ensures you **don't lose control of it.**"

CHAPTER 3

The ten pillars.

ISO/IEC 42001's requirements, translated into ten practical pillars for executive teams.

ISO/IEC 42001 organises its requirements technically. That structure works for auditors and certification bodies. It is less useful for boards. The ten pillars below are the same content, organised the way a senior leader would think about it — from strategic intent through to stakeholder trust.

Each pillar follows the same structure: what it is, what ISO/IEC 42001:2023 requires, what good looks like, and three questions to ask your organisation. Read them in any order. Return to whichever one is currently the most uncomfortable to discuss — that is usually where to start.

1. Strategic Intent

Why does AI exist in your business? What are you trying to achieve?

Strategic intent is the answer to the question, "why AI?" If your organisation cannot answer that in plain language, the rest of the governance framework has nothing to anchor to. AI without strategic intent becomes a series of disconnected experiments — some helpful, some wasted, none accountable. Strategic intent does not have to be sweeping. "We use AI to reduce manual work in finance and customer operations, in ways that are accurate, auditable, and consistent with our values" is a stronger statement of intent than "We are an AI-driven organisation."

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Alignment between AI use and the organisation's broader business strategy.
- Defined objectives for AI — what success looks like, in measurable terms.
- Identified use cases or domains where AI is being applied.
- A documented AI risk appetite — how much risk the organisation is prepared to take, in what domains, for what return.
- An ethical stance — what kinds of AI use are off the table for this organisation, regardless of return.

WHAT GOOD LOOKS LIKE

- AI use cases connect directly to a business objective the board can name.
- Success metrics exist and are reviewed.
- There is an explicit list of 'where we will not use AI', agreed at executive level.
- The strategic intent is written down in language a non-specialist can understand.
- The intent is reviewed at least annually as the AI landscape evolves.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- If a major customer asked us today, 'why are you using AI?', what would we say?
- Where in the business have we explicitly chosen not to use AI — and why?
- How would we know if our AI strategy was working?

THIS QUARTER

Write a one-paragraph statement of your AI strategic intent and test it against three named business objectives. Circulate it to the executive team and ask: does this reflect what we are actually doing?

For a detailed practitioner guide to Strategic Intent — including the AutoRide SA case study, use case prioritisation tools, and ROI framework — ask about our Pillar 1 Playbook.

2. Governance & Leadership

Who is accountable for AI in your organisation — and is that accountability real?

Governance only works when accountability is named. "IT is responsible for AI" is not accountability — it is delegation by default. Real governance answers the questions: Who decides? Who is accountable when something goes wrong? Who has the authority to stop something? Who briefs the board? In most organisations, AI accountability has been assumed by whoever happens to be using AI — a fragmented arrangement that holds up only as long as nothing goes wrong.

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Top management commitment to AI governance, documented and visible.
- An AI policy approved at executive or board level.

- Defined roles and responsibilities for AI — owners, custodians, users.
- Resources — people, budget, time — allocated to AI governance.
- Integration of AI governance with existing risk and assurance structures.

WHAT GOOD LOOKS LIKE

- There is a named senior executive accountable for AI governance — not a committee, a person.
- AI is on the board agenda at least quarterly, with substantive discussion.
- There is a clear policy that defines what AI use is acceptable, restricted, or prohibited.
- Decision-making authority for AI deployment is documented and understood.
- AI governance reports flow into existing risk and audit committees, not as a separate silo.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- Who is the named accountable person for AI in our organisation?
- When did the board last have a substantive conversation about AI — and what was the outcome?
- If an AI tool produced a bad outcome tomorrow, who would answer for it?

THIS QUARTER

Name the single executive accountable for AI governance in writing. Communicate this to the board. Add AI as a standing 15-minute item on the next board agenda.

For a detailed practitioner guide to Governance and Leadership — including governance architecture design, RACI matrices, and the AI Governance Health Check — ask about our Pillar 2 Playbook.

3. Risk & Impact Management

What could go wrong — and what are you doing about it?

AI introduces risks that traditional risk frameworks were not designed to capture. A spreadsheet error is a familiar risk. An AI model that systematically under-prices a particular customer segment is a different kind of risk — harder to detect, slower to surface, and capable of compounding before anyone notices. Risk and impact management for AI is not just a matter of asking 'what could go wrong?' It is asking 'what could go wrong in ways we wouldn't see?' — and putting controls in place to surface those failures.

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Systematic identification of AI-related risks — to people, to data, to reputation, to operations, to compliance.
- Impact assessments for high-stakes AI use cases, particularly those affecting individuals.
- Documented controls and mitigations for the risks identified.
- Ongoing risk evaluation as AI use evolves and as the AI landscape changes.
- Alignment with the organisation's wider enterprise risk framework.

WHAT GOOD LOOKS LIKE

- There is a register of AI-related risks, reviewed regularly, owned by named individuals.
- Risk assessments are conducted before new AI use cases are deployed, not after.
- Impact on affected stakeholders — customers, employees, third parties — is explicitly considered.
- Risk findings flow into the same risk reporting structure as other enterprise risks.
- There is a documented 'AI red list' — use cases or risk levels that require executive sign-off.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- What's on our AI risk register — and when did we last look at it?
- Which of our current AI uses would benefit from an impact assessment we haven't done?
- If our highest-stakes AI use case failed, what would the consequences be — and who would absorb them?

THIS QUARTER

Start an AI risk register. Even if it has only three entries, naming risks, owners, and mitigations is the foundation. Identify your highest-stakes AI use case and confirm whether a formal impact assessment has been conducted.

For a detailed practitioner guide to Risk and Impact Management — including risk taxonomy, impact assessment frameworks, and mitigation strategies — ask about our Pillar 3 Playbook.

4. AI Lifecycle Management

From the moment an AI use case is conceived to the moment it is retired — what controls govern its journey?

AI systems do not exist as point-in-time artefacts. They exist across a lifecycle: conception, design, validation, deployment, operation, maintenance, decommissioning. Each stage has its own governance requirements. A model that was carefully validated on launch can drift over time. A vendor that was responsible at procurement can be acquired by a different one a year later. Lifecycle management is the discipline of ensuring governance does not stop at deployment.

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Documented processes for each lifecycle stage: concept and use case definition, design and development, validation and testing, deployment, operation and monitoring, maintenance and improvement, decommissioning.
- Stage gates and approval points for moving an AI system from one stage to the next.
- Documentation, traceability, and version control for AI systems.
- Defined responsibilities at each stage — who does what, who approves what.
- A retirement plan for AI systems — how they will be safely shut down when their time comes.

WHAT GOOD LOOKS LIKE

- Every operational AI system can be traced through its lifecycle, with documentation at each stage.
- Validation and testing happen before deployment, not in production.
- Monitoring is in place for each operational AI system, with thresholds for intervention.
- There is a documented decommissioning process, used when systems are retired.
- The lifecycle is reviewed and improved based on what is learned from incidents and reviews.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- How many AI systems do we have in operation — and could we list them with confidence?
- When did we last formally validate a deployed AI system?
- What's our process for retiring an AI system we no longer use?

THIS QUARTER

Produce an inventory of AI systems currently in operation. Include vendor-embedded AI in tools you already use. Most organisations are surprised by the length of the list.

For a detailed practitioner guide to AI Lifecycle Management — including governance gate design, version control standards, and decommissioning protocols — ask about our Pillar 4 Playbook.

5. Data & Technology Governance

What runs your AI — and is it fit for purpose?

Every AI system depends on data. Garbage in, garbage out is the oldest principle in computing, and AI does not transcend it — it amplifies it. AI also depends on infrastructure: the platforms on which models run, the integrations through which they consume and produce data, the security controls that protect them. Data and technology governance ensures both layers are sound. It is the layer where the engineering discipline you already apply to other systems extends to the AI ones.

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Data quality, integrity, and provenance for the data feeding AI systems.
- Privacy and personal data protection aligned with applicable law (POPIA, GDPR, others).
- Data sourcing controls — where data comes from, including third-party data.
- Model transparency and explainability appropriate to the use case.
- Infrastructure and security controls for AI systems, equivalent to those applied to other critical systems.

WHAT GOOD LOOKS LIKE

- There is a clear inventory of data feeding AI systems, with quality and provenance documented.
- Privacy impact assessments are conducted for AI use cases involving personal data.
- Models are documented to a standard appropriate to their criticality.
- AI systems are subject to the same security controls as other operational systems.
- Third-party data and third-party AI services are governed with the same discipline as in-house ones.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- Where does the data feeding our AI systems come from — and how do we know it's reliable?
- If a regulator asked how a particular AI decision was made, could we explain it?
- Are our AI systems covered by the same security and access controls as our other critical systems?

THIS QUARTER

Identify the three AI systems in your organisation that process the most personal data. Confirm whether a POPIA-aligned privacy impact assessment has been conducted for each. If not, commission one.

For a detailed practitioner guide to Data and Technology Governance — including data quality frameworks, POPIA alignment, and model transparency standards — ask about our Pillar 5 Playbook.

6. Human Oversight

Where does a person stay in the loop — and is that oversight meaningful?

Human oversight is the most discussed and most often misunderstood aspect of AI governance. "Human in the loop" is easy to claim and difficult to make real. A person who rubber-stamps every AI output is not exercising oversight; they are providing legal cover. Real human oversight requires three things: the human has the information to make a meaningful judgement, the authority to override the AI, and the time to do so before it matters. Without all three, oversight is a ceremony.

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Defined human oversight requirements for each AI use case, calibrated to its risk level.
- Mechanisms for human intervention — escalation paths, override procedures, stop conditions.

- Decision accountability — the human (or role) who makes or confirms the final decision.
- Auditability and traceability of human decisions in AI-assisted workflows.
- Training and awareness so that humans in oversight roles understand what they are overseeing.

WHAT GOOD LOOKS LIKE

- For each material AI use case, the level and form of human oversight is documented.
- Humans in oversight roles have been trained on what to look for and when to intervene.
- Override decisions and their reasoning are logged and reviewable.
- Oversight is calibrated to risk — high-stakes uses get more, low-stakes uses get less.
- The oversight model is tested periodically — e.g. by reviewing whether interventions actually occur when they should.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- For our highest-stakes AI use case, what does 'human oversight' actually look like in practice?
- If a person disagreed with an AI output today, what would they do — and would they have the authority to act?
- How do we know our human oversight is meaningful, not theatrical?

THIS QUARTER

For your highest-stakes AI use case, document what human oversight actually looks like in practice: who reviews what, with what information, and with what authority to intervene. If you cannot document it, the oversight may not be real.

For a detailed practitioner guide to Human Oversight — including oversight calibration frameworks, intervention protocols, and auditability requirements — ask about our Pillar 6 Playbook.

7. Performance Evaluation

Is your AI working — and how would you know if it stopped?

Performance evaluation answers the question every board will eventually ask: are we getting the value we expected from AI, and is it operating within tolerance? Without measurement, you cannot answer either. AI systems can drift in ways that are invisible without monitoring — outputs degrade slowly, biases shift over time, edge cases that were rare become more common as use grows. Evaluation is the discipline of catching these things while they are still small.

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Defined performance indicators for AI systems — accuracy, fairness, reliability, business value.
- Ongoing monitoring of those indicators in operation.
- Drift detection — mechanisms for spotting when an AI system's behaviour changes over time.
- Internal audit and review processes for AI systems, on a defined cadence.
- Reporting of AI performance to executive and board level, in language they can engage with.

WHAT GOOD LOOKS LIKE

- Each material AI system has performance indicators, and they are monitored.
- Drift is detected before it becomes a problem, not after.
- AI performance is reported to the board as part of broader operational reporting.
- Internal audit treats AI systems as in-scope, not out-of-scope-because-too-technical.
- Findings from monitoring lead to action, not just reports.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- Can we name the performance indicators for our most important AI systems?
- How would we know if an AI system was drifting — and how long would it take us to know?
- When was the last time AI performance was discussed at the executive table?

THIS QUARTER

Define performance indicators for your most important AI system. If none exist, define three: one for accuracy or reliability, one for business value delivered, and one for stakeholder impact. Start tracking them.

For a detailed practitioner guide to Performance Evaluation — including the five-dimension AI value framework, ROI measurement methodology, and integrated dashboard design — ask about our Pillar 7 Playbook.

8. Continual Improvement

Are you learning from incidents, near-misses, and changing conditions — and is that learning getting back into the system?

AI is not a solved problem. The technology is changing, the regulatory environment is changing, the use cases are changing, and your organisation is changing. A management system that does not improve is a management system that decays. Continual improvement is the discipline that keeps your AI governance current. It also makes the difference between an organisation that has "written the policy" and one that is genuinely managing AI well.

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Mechanisms for capturing incidents, near-misses, and lessons learned.
- A process for translating those lessons into updated policies, controls, and training.
- Ongoing horizon-scanning for new risks, new technologies, new regulatory developments.
- Periodic review of the AI Management System itself, not just the AI systems it governs.
- Evidence of corrective and preventive action — not just identification of issues, but resolution of them.

WHAT GOOD LOOKS LIKE

- Incidents and near-misses are logged and reviewed.
- There is a regular cadence for reviewing and updating the AI policy and AI Management System.
- Lessons from outside the organisation — industry incidents, regulatory developments — are captured.
- Issues identified in audits and reviews lead to documented changes within defined timeframes.
- The improvement loop is itself audited — someone is checking whether what was supposed to change actually did.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- When was the last AI-related incident or near-miss in our organisation — and what changed as a result?
- How do we keep our AI policy current with a fast-changing technology?
- Are we learning from other people's AI incidents — or only our own?

THIS QUARTER

Review the last three AI-related incidents or near-misses in your organisation. For each, identify what changed as a result. If nothing changed, that is your starting point for a continual improvement process.

For a detailed practitioner guide to Continual Improvement — including incident management frameworks, horizon-scanning processes, and management system review protocols — ask about our Pillar 8 Playbook.

9. Compliance & Assurance

Can you demonstrate to a regulator, an auditor, or a board that your AI governance is real?

Compliance is not the goal of AI governance — good outcomes for stakeholders is the goal. But compliance is the test. If you cannot demonstrate, with evidence, that your governance is operating as designed, then you do not have governance — you have intentions. Assurance is the discipline of producing that evidence and submitting it for examination, internally and externally, on a regular basis.

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Compliance with applicable laws and regulations — POPIA, sector-specific rules, international frameworks where relevant.
- Alignment with ISO/IEC 42001 (or equivalent), and where relevant, formal certification.
- Internal audit coverage of the AI Management System, conducted by appropriately skilled auditors.
- Documentation that supports examination by external parties — customers, auditors, regulators.
- An assurance map showing how different forms of assurance combine to give confidence.

WHAT GOOD LOOKS LIKE

- There is a clear list of laws, regulations, and standards the AI Management System is designed to comply with.
- Internal audit covers AI on a defined cycle.
- Documentation is maintained at a level that supports external examination.
- Assurance gaps — areas where neither internal nor external assurance applies — are identified and addressed.
- Where ISO/IEC 42001:2023 certification is in scope, the path to certification is documented and resourced.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- If an external auditor arrived tomorrow asking about our AI governance, what would we hand them?
- Where are the gaps between what we say we do and what we can demonstrate we do?
- Is ISO/IEC 42001:2023 certification a useful path for us — and if so, when?

THIS QUARTER

Map your current AI governance against the seven core areas of ISO/IEC 42001:2023. Identify the two areas with the most significant gaps. Commission a focused assessment of those areas before attempting to address the full standard.

For a detailed practitioner guide to Compliance and Assurance — including ISO/IEC 42001:2023 gap assessment tools, internal audit frameworks, and certification pathway guidance — ask about our Pillar 9 Playbook.

10. Stakeholder Trust & Transparency

How do you communicate about your AI use — and would your stakeholders trust the answer?

Trust is a slow currency. It takes years to build and minutes to lose. AI is currently producing more lost-trust events than any other technology category, because the technology is moving faster than the communication discipline around it. Stakeholder trust and transparency is the work of being clear with the people affected by your AI use about how it is being used, what choices they have, and what protections exist. It is also the work of being honest when something goes wrong.

THE ISO/IEC 42001 LENS — WHAT THIS PILLAR REQUIRES

- Transparency about where AI is in use, particularly in stakeholder-facing applications.
- Disclosure mechanisms appropriate to the audience — customers, employees, regulators, partners.
- User feedback mechanisms — ways for affected stakeholders to raise concerns.
- Clear communication about how AI decisions are made, particularly those affecting individuals.
- Trust frameworks — the cumulative evidence base that demonstrates the organisation's AI use is responsible.

WHAT GOOD LOOKS LIKE

- Customers know when they are interacting with AI — it is not concealed.
- Employees know how AI affects their work and have channels to raise concerns.
- There is a clear external statement of the organisation's AI principles, available to anyone who asks.
- Disclosures are written for the audience — customers get customer language, regulators get regulator language.
- The organisation has a plan for what to communicate, and how, when an AI incident occurs.

THREE QUESTIONS TO ASK YOUR ORGANISATION

- If a customer asked us today how we use AI in our service to them, could we answer clearly?
- Do our employees know how AI affects their work — and trust how we are managing it?
- If something went seriously wrong with one of our AI systems, what would we say, to whom, and how quickly?

THIS QUARTER

Draft a one-page AI transparency statement for your organisation: what AI you use, why, and what protections exist for those affected. Test it with one customer and one employee. The gaps in their understanding are your communication priorities.

For a detailed practitioner guide to Stakeholder Trust and Transparency — including disclosure frameworks, trust measurement tools, and crisis communication protocols — ask about our Pillar 10 Playbook.

Where are you now? A ten-pillar self-assessment.

Use this assessment to take stock of where your organisation stands across the ten pillars. For each pillar, mark your current position honestly. This is not a test. It is a map. The pillars where you mark Not in place or Partially in place are where to focus first.

Pillar	In place	Partially in place	Not in place
1. Strategic Intent	☐	☐	☐
2. Governance & Leadership	☐	☐	☐
3. Risk & Impact Management	☐	☐	☐
4. AI Lifecycle Management	☐	☐	☐
5. Data & Technology Governance	☐	☐	☐
6. Human Oversight	☐	☐	☐
7. Performance Evaluation	☐	☐	☐
8. Continual Improvement	☐	☐	☐
9. Compliance & Assurance	☐	☐	☐
10. Stakeholder Trust & Transparency	☐	☐	☐

Three or more pillars Not in place: start with Strategic Intent and Governance. These are the foundations everything else builds on. Three or more Partially in place: the AI Opportunity and Risk Diagnostic is the right next step. It will identify which partial implementations are sound and which need rebuilding. All ten In place: ISO/IEC 42001:2023 certification may be achievable in 12 to 18 months. Ask us about the readiness pathway.

For a facilitated version of this assessment with benchmarking against South African peers, take the AI Fluency Assessment at https://theaifluencyco.com or book a 30-minute conversation at https://calendly.com/zoefrancois123/30min.

ABOUT

The AI Fluency Company.

The AI Fluency Company helps executive teams adopt AI responsibly — building the language, the judgement, and the governance to make AI work in their business. We work with mid-market organisations, family businesses, and compliance-minded leaders who want to do AI properly the first time.

How we work with executive teams.

ENGAGEMENT	WHAT IT IS
AI Fluency Training	Half-day workshop for leadership teams. Builds shared understanding of working with AI effectively, efficiently, ethically, and safely.
AI Strategy Workshop	Two days. Prioritise use cases, map ownership, build a 90-day roadmap.
AI Opportunity & Risk Diagnostic	Five days. Board-ready review of where AI belongs in your business and where you are exposed. ISO/IEC 42001 aligned.
Hyperautomation Programme Design	Structured engagement to apply this playbook to a specific transformation opportunity — from discovery through deployment readiness.
AI Governance & ISO 42001 Readiness	A structured pathway to certification, delivered with an accredited specialist partner.

About the author.

Zoë François is the founder of The AI Fluency Company. She trained as a chemical engineer and built one of South Africa's earliest applied industrial AI systems in the 1990s — an expert system on the Gensym G2 platform that modelled Anglo American Platinum's refinery operations. She then spent two further decades in commercial leadership, ultimately as **General Manager of Sales at Sappi** and **Senior Sales Executive at Consol Glass**. She holds the Oxford Saïd Business School AI Programme certificate.

Engineering rigour. Commercial judgement. The combination most AI engagements are missing.

TAKE THE NEXT STEP

Take the AI Fluency Assessment for a maturity benchmark : https://theaifluencyco.com

Book a 30-minute conversation to walk through your AI governance starting point. No charge for the first call.

Subscribe to receive future editions of this guide and other thinking on AI governance for executive teams.

First edition · 2026 · Cape Town