



When the Board Asks About Your AI: What You Need to Have Ready

AI governance has moved from IT policy to board agenda. What auditability, transparency, and responsible AI mean in practice, and what a POPIA-compliant AI deployment actually looks like.

By **Mike Walker** · Founder, SL8 Agency

The shift happened quietly. Twelve months ago, AI governance was a compliance checkbox owned by IT and legal. Today it's on the board agenda. Boards are asking questions IT teams aren't always ready to answer.

Here are the questions, and what you need ready before they land.

"What decisions is our AI making?"

Not what the system is capable of. What decisions it is actually making, in production, today. Most organisations can't answer that with any precision. They know the system exists. They know roughly what it does. They can't produce a log of decisions made, why they were made, and what data went in.

That's an audit trail problem. Any AI system you build for production needs to log its decision logic. Not to satisfy a regulator today, but because the regulator will ask tomorrow.

"Who is responsible when it goes wrong?"

The law hasn't settled this yet. Boards are asking anyway. The practical answer has to be a named person with a defined role. "The AI made the decision" doesn't survive contact with a board, a customer, or a journalist.

Before you deploy an AI system into a customer-facing process, you need a defined escalation path: what happens when it produces a wrong output, who catches it, and who owns the remediation.

"Is our customer data safe?"

POPIA is the local frame. The relevant questions: where is the data going when a user prompt is processed? Is it leaving South Africa? Is it being used to train a third-party model? Is it stored, and if so, where and for how long?

If you're routing customer data through a US-based SaaS AI provider on default retention settings, the answer to most of those questions is one you don't want to give a board.

The cleanest POPIA answer is a private deployment: model runs in your environment, data never leaves your tenancy, you hold the audit trail. That isn't practical for every use case, but it should be the starting position for any system touching customer data.

"What is the ROI?"

This one has always been there, but boards are asking it with less patience. "We're exploring AI to stay competitive" stopped being an answer about a year ago. The answer needs a number: hours saved, error rate reduced, revenue influenced, cost removed.

If you can't attach a number to the system, the board will defund it eventually. Not because they don't believe in AI. Because everything else in the budget has a number.

What a board-ready AI deployment looks like

Four things:

- 1 An audit trail of decisions: logged, queryable, retained
- 2 Named accountability: a person, not a department
- 3 Data residency documentation: where data goes, when, and why
- 4 A business case with a number: before go-live, not after

None of these require advanced technical capability. They require process discipline applied at the point of deployment.

The organisations that will scale AI infrastructure through the next board cycle are the ones treating governance as an engineering requirement rather than a post-launch concern.

Mike Walker

Founder, SL8 Agency

Ready to stop experimenting and start building?

One conversation. A straight read on where you actually stand.

BOOK A CALL



If the button doesn't open: `calendar.app.google/7J2ThQMgXULCqeZe6`