
The Enterprise AI Playbook

From LLMs to AI Co-Workers

A practical guide to RAG, AI agents, automation, and governance.

Kennedy Maweu

Founder of Enterprise AI Africa

WHY THIS BOOK MATTERS

AI activity is not AI value.

Most AI advice was written for a different market. It assumes cheap cloud credits, clean data, large IT teams and customers who buy online. In Nairobi, Lagos or Kampala the reality is different. Data sits in spreadsheets, WhatsApp threads and legacy systems. IT is a team of three. Staff move between English and Swahili in the same sentence. Regulators expect evidence, not enthusiasm.

So organisations experiment. They buy licences, run a workshop and celebrate a pilot. Months later the board asks one question: what did it change in the P&L? Nobody can answer.

AI activity is not a proxy for AI value. The gap between the two is a design problem, not a technology problem. This book shows how to close it: choose the right workflows, give AI the ability to act inside them, govern it, and measure the result.

THE BOARD TEST

Can you point to a line on the P&L that your AI changed?

What is inside the **full book**

01 **Why AI Activity Is Not AI Value**

The gap between spend, pilots and the P&L

02 **From Chatbots to AI Co-Workers** **IN THIS SAMPLE**

How an agent differs from a writing assistant

03 **Giving AI Your Company's Knowledge**

RAG explained for executives, without the hype

04 **Designing Agents That Take Action**

Goals, tools, permissions and human approval

05 **Workflows First**

Choosing the processes worth automating

06 **Governance and Guardrails**

Policy, data boundaries and decision rights under Kenya's Data Protection Act 2019

07 **Measuring ROI**

From pilot to P&L: the numbers a CFO will accept

08 **Your 90-Day Deployment Plan**

From first workflow to governed, measurable adoption

CHAPTER 2

From Chatbots to AI Co-Workers

A standard large language model (LLM) is a writing assistant. You give it text, it gives text back. It can draft an email, summarise a contract or rewrite a proposal, and it does this well. What it does not do is anything else. It does not know your pipeline, it cannot open your CRM, and it forgets the task when the conversation ends. Every useful output depends on a person remembering to ask, then copying the result somewhere else.

An AI agent is a different kind of system. It uses an LLM as its reasoning layer, but it is connected to your data and your tools, and it is given a goal, not a prompt. It can look up records, decide what to do next, take the action and report back. The shift is from generating text to completing work.

Consider a sales team of eight people.

Version one: the writing assistant. Each rep opens a chat window and asks for a better cold email. The emails improve. But who to contact, in what order and when to follow up still depends on each rep's memory and mood. Leads go cold on Friday. The CRM is updated on Monday, if at all. The team is faster at the same manual process, and the pipeline barely moves.

Version two: the AI co-worker. An agent reads new leads as they enter the CRM. It checks each one against criteria you defined: company size, sector, past engagement, signals of intent. It scores and ranks them. It drafts the first message for the top prospects, schedules follow-ups at set intervals, logs every touch back to the CRM, and sends the team a ranked list each morning. A person reviews and approves anything that reaches a client. Nothing slips, because nothing depends on someone remembering.

The difference is not intelligence. It is architecture. The same model can power both versions. What changes is what it is connected to, what it is allowed to do, and who checks its work.

Why this matters for ROI

A writing assistant saves minutes per task, and those minutes are hard to find on a P&L. An agent changes a process: more leads contacted on time, fewer deals stalled, a pipeline you can see. Those are numbers a CFO can track.

It also raises the stakes. A system that acts can act wrongly. Before any agent goes live, define three things: what data it can access, what actions it may take without approval, and how its outcomes are measured. Skip them and you have built a faster way to make mistakes.

The question for your next board meeting is not "which AI tool should we buy?" It is "which workflow should an AI co-worker own, and how will we know it worked?"

Ready to build your **AI Playbook**?

This sample covers one chapter. The full book gives your leadership team the working tools to move from experiments to governed, measurable AI:

- **The frameworks.** Step-by-step methods to choose workflows, design agents and set approval rules.
- **The governance checklists.** Data boundaries, decision rights and risk controls you can adopt and adapt.
- **The case studies.** Worked examples that show how the numbers change when AI owns a process.
- **The 90-day plan.** A deployment path your CFO and COO can sign off.

BUY THE FULL EBOOK - KES 4,500

enterpriseaiafrica.africa/book-registration



ABOUT THE AUTHOR

Kennedy Maweu is the Founder of Enterprise AI Africa and a Claude Certified AI Architect (Anthropic), with AI Security and Governance certification. He builds governed AI systems for East African organisations. Learn more at enterpriseaiafrica.africa.