AI acceptable-use policy

An AI acceptable-use policy your teams will actually follow

From the practitioner
Wanjiru Kamau
Former Head of Data and Analytics, an East African bank

PDF manual · Reviewed November 2026 · portsandbridges.com

OPERATING MANUAL · AI ADOPTION

An AI acceptable-use policy your teams will actually follow

A manual for putting an AI acceptable-use policy in place in a company of one to five hundred people. It covers which tools are approved, what data must never be entered, how output is checked before it is used, and who is accountable, then shows how to introduce the policy so teams adopt it rather than work around it.

For: CEOs, COOs, heads of risk and technology · Time: Two to four weeks to agree and introduce

When to use this manual

Use this manual when people in your company are already using AI tools, whether or not anyone has approved them, or when leadership wants to encourage wider use safely. It suits companies regulated under the Nigeria Data Protection Act, the Kenya Data Protection Act, Ghana's Data Protection Act or similar laws, where personal data in the wrong tool is a reportable breach.

Before you start

[ ] A named owner for the policy with time to run it, usually in risk, technology or operations
[ ] A list of the AI tools people already use, gathered by asking rather than auditing
[ ] Your data classification, or agreement to use the simple three-level version in this manual
[ ] Sponsorship from the CEO, so the policy is seen as enabling work rather than blocking it

The sequence

1. Name the owner and the purpose

Make one person accountable for the policy and its upkeep.

Decision	Who owns the policy, and whether its aim is to enable use, limit risk, or both
Owner	CEO
Check before moving on	The owner and purpose announced to the leadership team

2. Approve the tools

Give people safe, approved options so they stop using personal accounts.

Decision	Which tools are approved, on which plans, and for which teams
Owner	Policy owner with technology
Check before moving on	Each approved tool's terms reviewed for data retention and training on your data

3. Classify the data

Make it obvious what can and cannot go into an AI tool.

Decision	What counts as public, internal and restricted data, with examples from your business
Owner	Policy owner with the data protection officer
Check before moving on	Ten real examples classified by a mixed group of staff, with the same answers

4. Set the review rules

Make sure AI output is checked before it reaches a customer, regulator or decision.

Decision	Which uses need human review, by whom, and what they check
Owner	Function heads
Check before moving on	Review rules written into the process for each approved use

5. Define accountability

Make clear that people remain responsible for work they produce with AI.

Decision	What happens when the policy is breached, and how incidents are reported
Owner	Policy owner with HR
Check before moving on	Incident reporting route tested with a sample case

6. Introduce and train

Get the policy used, not filed.

Decision	How the policy is introduced, and who must complete training
Owner	Policy owner with managers
Check before moving on	Every manager has walked their team through the policy with their own examples

7. Review every quarter

Keep the policy current as tools and uses change.

Decision	Which tools to add or remove, and which rules to change
Owner	Policy owner
Check before moving on	A quarterly review held, with changes published to staff

Who decides what

Decision	Leadership	Policy owner	Managers
Purpose of the policy	Approves	Proposes	Informed
Approved tools	Informed	Owns	Requests
Data classification	Approves	Owns	Applies
Review rules for a use	Informed	Sets the minimum	Owns for their team
Handling a breach	Informed of serious cases	Owns	Reports

Where it goes wrong

A policy that only says no

Early sign: Usage goes quiet officially while personal accounts keep appearing
What to do: Approve at least one tool on day one and lead with what people can do

Rules nobody can apply

Early sign: Staff ask the policy owner whether every task is allowed
What to do: Use three data levels with examples from your own business, not abstract definitions

No owner after launch

Early sign: The approved tool list is six months old and new tools are in use
What to do: Put the quarterly review in the owner's objectives and the risk committee's calendar

The checklist

[ ] A named owner for the policy with time to run it, usually in risk, technology or operations
[ ] A list of the AI tools people already use, gathered by asking rather than auditing
[ ] Your data classification, or agreement to use the simple three-level version in this manual
[ ] Sponsorship from the CEO, so the policy is seen as enabling work rather than blocking it
[ ] The owner and purpose announced to the leadership team
[ ] Each approved tool's terms reviewed for data retention and training on your data
[ ] Ten real examples classified by a mixed group of staff, with the same answers
[ ] Review rules written into the process for each approved use
[ ] Incident reporting route tested with a sample case
[ ] Every manager has walked their team through the policy with their own examples
[ ] A quarterly review held, with changes published to staff

From the practitioner

> "Banning AI does not stop people using it; it only stops them telling you. The policies that work are short, name the tools people may use, and are clear about the handful of things that must never happen, such as customer data going into a public tool. Write it for the analyst on a deadline, not for the auditor."

Wanjiru Kamau, Former Head of Data and Analytics, an East African bank

Questions

What should an AI acceptable-use policy include?

The purpose of the policy, the approved tools, what data may and may not be entered, when output must be checked by a person, who is accountable and how to report problems. Keep it to a few pages so people read it.

Can staff put customer data into ChatGPT or similar tools?

Not into public or personal versions. Personal data entered there may be stored and used outside your control, which can breach data protection law. Enterprise plans with the right terms can be approved for some uses after review.

Who should own the AI policy?

One named person with time to run it, usually in risk, technology or operations, supported by the data protection officer. Shared ownership tends to mean nobody updates it.

How often should the policy be reviewed?

Every quarter while tools change quickly, and whenever a new tool is approved or an incident occurs.

Want a practitioner to help you put this to work? Start a diagnostic at portsandbridges.com/diagnostic or write to info@portsandbridges.com.