WIPHLA LEADERSHIP SESSION · 13 NOVEMBER 2026

AI Use Policy Starter Kit

Everything from Anatomy of an AI Use Policy, in the order you will need it. Start on page one on Monday morning.

The one idea to hold. Wherever AI shapes a decision about people, money or trust, a named person owns it, can see enough to question it, and can stop it.

The Judgment Line™ is the boundary in an AI-supported process where accountable human authority must become explicit.

Prepared by **Dr. Aday E. Adetosoye** **Free, courtesy of WiPHLA**

INSIDE THIS KIT

01 Your first 30 days

02 The ten questions, with sample wording

03 Three zones and the never list

04 Before you hit send: five checks

05 Interim guidance, ready to send

06 Four-question staff survey

07 Before you approve an AI tool

08 If something goes wrong

09 Amara's story, worked through

10 Notes and sources

01

Your first 30 days

You can start this week with a named owner, a few clear ground rules and the plan below. In about 30 days you will have a first draft, ready for your next board meeting. Bring it to the WiPHLA Policy Clinic on Friday, 11 December, to strengthen it with peers. If you advise a leader rather than lead, start with three moves: forward the interim guidance to your CEO, offer to run the staff survey, and use the five checks on your next analysis. Tick items off as you go; your ticks stay in this browser.

WEEK 1 Find out where you stand

- [ ] Confirm who owns AI policy. If nobody does, propose a named owner to your CEO.
- [ ] Send the four-question staff survey (section 06).
- [ ] List every AI tool in use, including free apps on personal phones and AI features inside software you already pay for.
- [ ] Name the three decisions where AI already shapes outcomes for people, money or trust.

WEEK 2 Set ground rules now

- [ ] Adapt the interim guidance (section 05) and send it to all staff, volunteers and consultants.
- [ ] Set up one place for questions and concerns: a named person or a shared email address.
- [ ] Ask your data protection lead or legal adviser which national data laws apply to your work.

WEEK 3 Draft the policy

- [ ] Answer the ten questions (section 02), starting from the sample wording.
- [ ] For each of your three AI-shaped decisions, name a decision owner, a reviewer and a deputy.
- [ ] Run every tool in use through the tool approval checklist (section 07).
- [ ] If you have a policy buddy, swap first drafts by Friday, 4 December, and trade one honest round of feedback.

WEEK 4 Test and approve

- [ ] Walk one real workflow through the draft from start to finish, and fix any gaps you find.
- [ ] Take the draft to leadership, and put it on the agenda for your next board meeting.
- [ ] Bring your draft to the WiPHLA Policy Clinic on Friday, 11 December.
- [ ] Schedule role-based training and put the first review date on the calendar, six months out.

02

The ten questions your policy must answer

Your policy is a set of answers. Work through these in order, adapt the sample wording to your organization and country, and fill every highlighted placeholder.

GROUND RULES

1 Why are we using AI?

Decide the value you expect, and what you must protect.

> *"We use AI to save staff time and improve services, while protecting the safety, privacy, rights and trust of the communities we serve."*

Done when any staff member can explain in one sentence why the organization uses AI.

2 Who and what does it cover?

Decide the people, devices, partners, and what counts as AI.

> *"This policy covers staff, volunteers, consultants and board members using AI for our work on any device, including personal phones and messaging apps. Partners working on our behalf follow it through their agreements. AI includes tools that generate text, images or audio, and tools that rank, score or predict."*

Done when consultant and partner agreements refer to the policy.

3 What can staff use AI for?

Decide which uses sit in each zone: Delegate, Verify or Own (section 03).

> *"Staff may use approved AI tools for authorized work. Every use has an owner who checks the output in proportion to the stakes. A qualified person verifies any*

content others will rely on. A named person makes and records every decision about people, money or trust."

Done when the policy lists at least two real examples of your own for each zone.

---

4 What is never allowed?

Decide your prohibited uses. Start with the five in section 03.

"The following uses are prohibited under all circumstances: [list]."

Done when every staff member has seen the list.

PROTECTIONS

---

5 Which tools are approved?

Decide who approves tools and uses, and on what evidence.

"The [operations lead] approves each AI tool, and each use of it, before work data enters it. Approval follows a documented check of data storage and transfers, use of our data for training, security, performance in the languages we work in, contract terms, and how we would exit."

Done when a current list of approved tools and approved uses exists and staff can find it.

---

6 What data can go in?

Decide what data each approved use may receive.

"Before personal or health data enters an AI tool, the [data protection lead] confirms the legal basis under applicable national law, limits the data to what the task requires, confirms a signed data agreement with the vendor, and checks any processing outside the country. Identifying information never enters unapproved tools."

Done when staff know which data may go into which tool.

7 Who decides, and who checks? THE JUDGMENT LINE

Decide a decision owner, a reviewer and a deputy for each high-stakes use.

> *"For each AI-supported decision about people, money or trust, a named decision owner makes the decision and records the reasons. A qualified reviewer sees the full evidence, including what the tool excluded, and may change, escalate or pause the decision. Where a board or committee holds formal authority, a named person prepares and owns the recommendation that reaches it."*

Done when every high-stakes use has three names (owner, reviewer and deputy) and the reviewer passes the See, Change and Stop tests.

PRACTICE

8 Who do we tell, and how can they challenge?

Decide whom you tell, how, and the route for raising concerns.

> *"When AI materially shapes a decision that affects a community, partner, funder or person, we tell them in plain language. We offer a reachable person, in their language and by phone or in person where needed, who can explain the decision and correct errors."*

Done when a challenge route exists and someone has tested it.

9 What happens when something goes wrong?

Decide your incident lead and first steps (section 08).

> *"Anyone may report an AI concern without blame. The [incident lead] pauses unsafe use, assesses which data, people and decisions were affected, makes*

corrections, meets any legal notification duties, and approves any restart."

Done when staff know whom to call, and the incident lead has walked through one practice scenario.

10

How do we keep it current?

Decide training, review rhythm, and what triggers a fresh review.

> *"Staff complete role-based training before access. The policy owner reviews every approved use every six months and after any change in tool, data, purpose, scale or underlying data system, and renews, restricts or retires each use."*

Done when the next review date is on the calendar.

03

Three zones and the never list

Every use has an owner. Checking grows with the stakes.

Delegate

WHEN

Low-stakes work using public or non-sensitive information

EXAMPLE

A draft meeting agenda

CHECK

You read it fully before you use it

Verify

WHEN

Content others will rely on

EXAMPLE

A donor report, or a ranking prepared for the board

CHECK

A qualified person checks every fact and figure against the source

Own

WHEN

Decisions about people, money or trust

EXAMPLE

Which districts receive funds

CHECK

A named person decides with full evidence, and can stop it

Never allowed, whatever the checks

1. Putting names, photos, case notes or other identifying details into tools that are not approved for them
2. Letting AI alone decide anything about a person's care, eligibility, job or safety
3. Sharing AI output as fact before a person has checked it
4. Creating images, quotes or stories that pass as real people or beneficiaries
5. Putting internal or confidential information into unapproved tools

04

Before you hit send: five checks

For anyone preparing an AI-assisted analysis that others will act on. These are the checks Amara needed before her ranking went to the board.

1 Check the data's age	Add a "last reported" column. Flag every unit with stale or missing records, especially after a system migration.
2 Ask how the tool ranked	Which criteria did it use, and what did it do with missing data?
3 Recheck against the source	Rebuild the top and bottom of the list yourself from the original data.
4 Test it against the field	Ask two people closest to the work whether the result matches what they see.
5 Show the full picture	Share the full list with gaps flagged, a one-line method note, and a recommendation for anything uncertain.

The person who prepares the analysis owns its accuracy. The decision owner owns the decision. Write both roles into your policy so neither falls through the cracks.

05

Interim guidance, ready to send

Send this in week 2. It protects people while you draft the full policy. Copy it, then replace every highlighted placeholder.

Copy the guidance

Subject: Using AI at `[Organization]` **: interim guidance**

Dear colleagues,

Many of us already use AI tools in our work. Until our full AI use policy is approved, please follow these six rules for any work use of AI, on any device, including personal phones.

1. **Use approved tools only.** For work, use only: `[list of tools]`. Ask `[name]` before using any other.
2. **Protect people's information.** Never put names, photos, case notes or anything that could identify a person or community into an AI tool unless `[name]` has approved that tool for that data.
3. **You own what you produce.** Read anything AI helps you produce before you use it. Check every fact and figure that others will rely on against the original source.
4. **A person decides.** AI may inform decisions about someone's care, eligibility, job, funding or safety. A named person makes those decisions.
5. **Be open about it.** Tell your manager when AI helped produce anything going to a funder, ministry, partner or community.
6. **Speak up.** If something goes wrong, or you are unsure, contact `[name]` at `[contact]`. Raising a concern will never be held against you.

Thank you for helping us use AI safely and well.

`[Name]`
`[Title]`

06

Four-question staff survey

Send it in week 1. Keep it anonymous so people answer honestly; the goal is to learn, and no answer leads to blame.

1. Which AI tools do you use for work, if any?
   Include free apps, phone apps, and AI features inside software we already use.

2. What do you use them for?

3. What kind of information have you put into them?
   Public information · Internal documents · Information about individual people or communities · Not sure

4. What would help you use AI well and safely at work?

07

Before you approve an AI tool

Approve the use as well as the tool. A tool approved for drafting reports needs a separate approval before it ranks, screens or decides.

- [ ] Where is our data stored, and does it leave the country?
- [ ] Does the vendor use our data to train its AI, and can we switch that off?
- [ ] How long does the vendor keep what we put in and what it produces?
- [ ] Who at the vendor, and which subcontractors, can access our data?
- [ ] Have we tested it on our own material, in the languages and settings we work in?
- [ ] Does the contract cover security, breach notification and deleting our data?
- [ ] Is the specific use approved, with an owner named?

- [ ] How would we stop using it, and get our data back?

Free and paid versions of the same tool can treat your data very differently, so check the specific plan, settings and contract.

08

If something goes wrong: the first 24 hours

- [ ] **Stop.** Pause the tool or the use, if it is safe to do so.
- [ ] **Report.** Tell the incident lead. Describe what happened without copying sensitive data into the report.
- [ ] **Assess.** Identify which data, which people and which decisions were affected.
- [ ] **Correct.** Fix wrong outputs, and tell everyone who received them.
- [ ] **Notify.** Check legal and contract duties to inform regulators, funders or affected people.
- [ ] **Learn.** Change the rule or process that allowed it, and record what changed.

Deleting a chat may leave copies a vendor already holds, so treat data exposure as real even after deletion.

09

Amara’s story, worked through all ten questions

Amara is the M&E manager. With the board pack due Thursday, she asked an AI tool to rank 40 districts by need for limited outreach funds, formatted its top 10, and sent the list up. The board approved it.

Four months later, a field officer asked why a district with a cholera outbreak received nothing. A DHIS2 migration had dropped three districts' recent records, and the tool read the gaps as low need. With the policy in place, each question changes the outcome.

QUESTION	WHAT THE POLICY REQUIRES IN AMARA'S CASE
1 · Why	Fair reach to high-need districts counts as part of success.
2 · Who and what	The policy covers Amara, the program team, and any consultant who runs a ranking.
3 · Which zone	Amara's analysis is Verify work. Allocating the funds is Own: a named person decides.
4 · Never	The tool may not make the final selection alone.
5 · Tools	The tool needs approval for ranking, separate from any approval for drafting.
6 · Data	Amara checks every district's data for age and gaps before the ranking runs, with extra care after the DHIS2 migration.
7 · Who decides	Amara owns the accuracy of the analysis. The program director owns the allocation and records the reasons. A reviewer sees all 40 districts and can pause.

QUESTION	WHAT THE POLICY REQUIRES IN AMARA'S CASE
8 · Tell and challenge	Districts learn how selections were made and can ask a named person for a review.
9 · When it goes wrong	When the cholera district surfaces, the team corrects the allocation, informs the board, and checks the other migrated records.
10 · Keep current	Any system migration, new data source, weighting or region triggers a fresh review.

10

Notes and sources

This kit is a practical starting point. Adapt it to your country's laws, your sector's rules and your organization's governance, and confirm data protection requirements with a qualified adviser. All case examples are fictional and describe no real organization.

Why now. On 22 October 2026 in Addis Ababa, African Union Member States validated the African Continental Health Data Governance Framework. Africa CDC will finalize it, with consideration through AU policy processes expected in 2027. Organizations that set their own AI and data rules now will be ready to align. (Africa CDC)

Further reading

- African Union Continental Artificial Intelligence Strategy (2024)
- WHO: Ethics and governance of artificial intelligence for health (2021)
- WHO: Guidance on large multi-modal models (2024)

Dr. Aday E. Adetosoye

The Judgment Line™ is a framework of Dr. Aday E. Adetosoye. © 2026 1099 Impact Street. This kit is free, courtesy of WiPHLA.

Questions: aday@1099impactstreet.com

The crown is lighter when we carry it together.