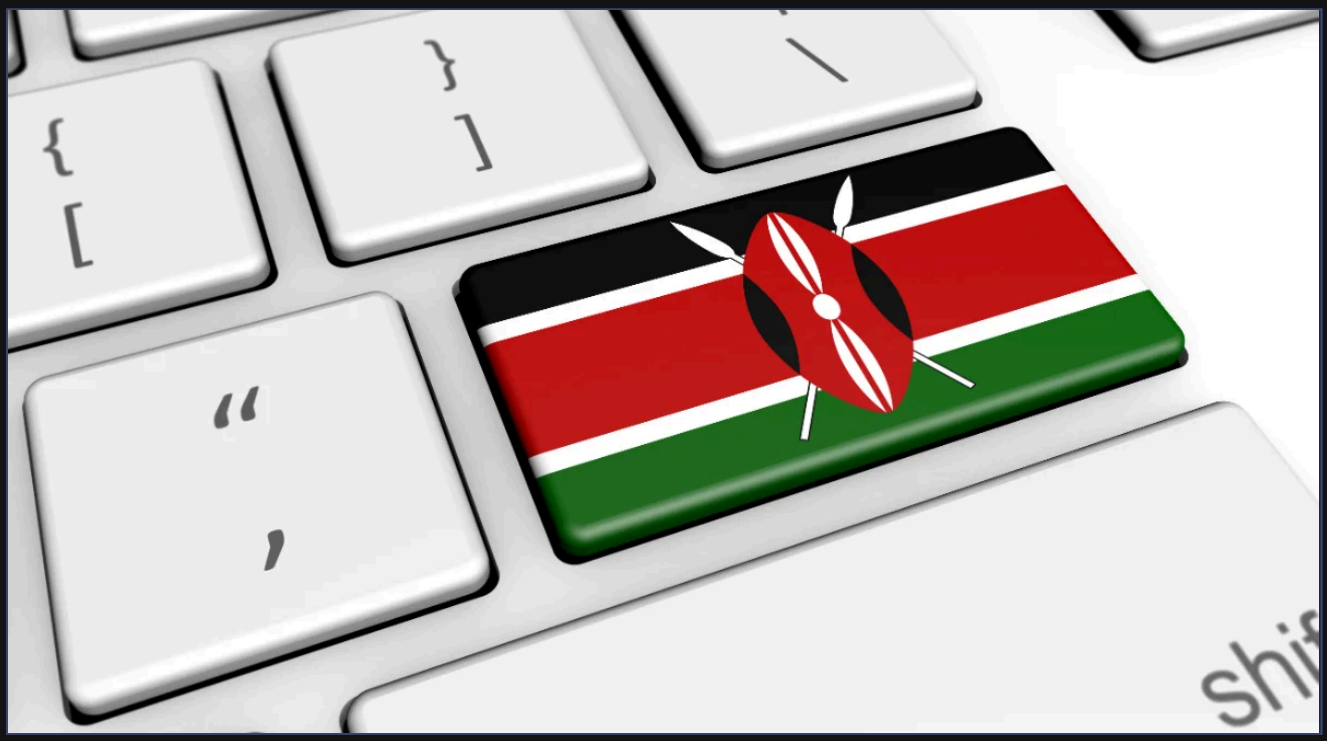
AI READINESS IN KENYA 2026 · WHITE PAPER

AI is running ahead of the rules.

What 214 Kenyan organisations told us about AI adoption, governance and risk, and what business leaders should do in the next 90 days.

November 2026

Findings from the KNCCI Africa Commerce & Industry Summit and the Safaricom Cybersecurity Summit



AKILI AI

60 Years of ADVANCING TRADE & INDUSTRY | KNCCI KENYA NATIONAL CHAMBER OF COMMERCE & INDUSTRY — Growing your Business together

CONTENTS

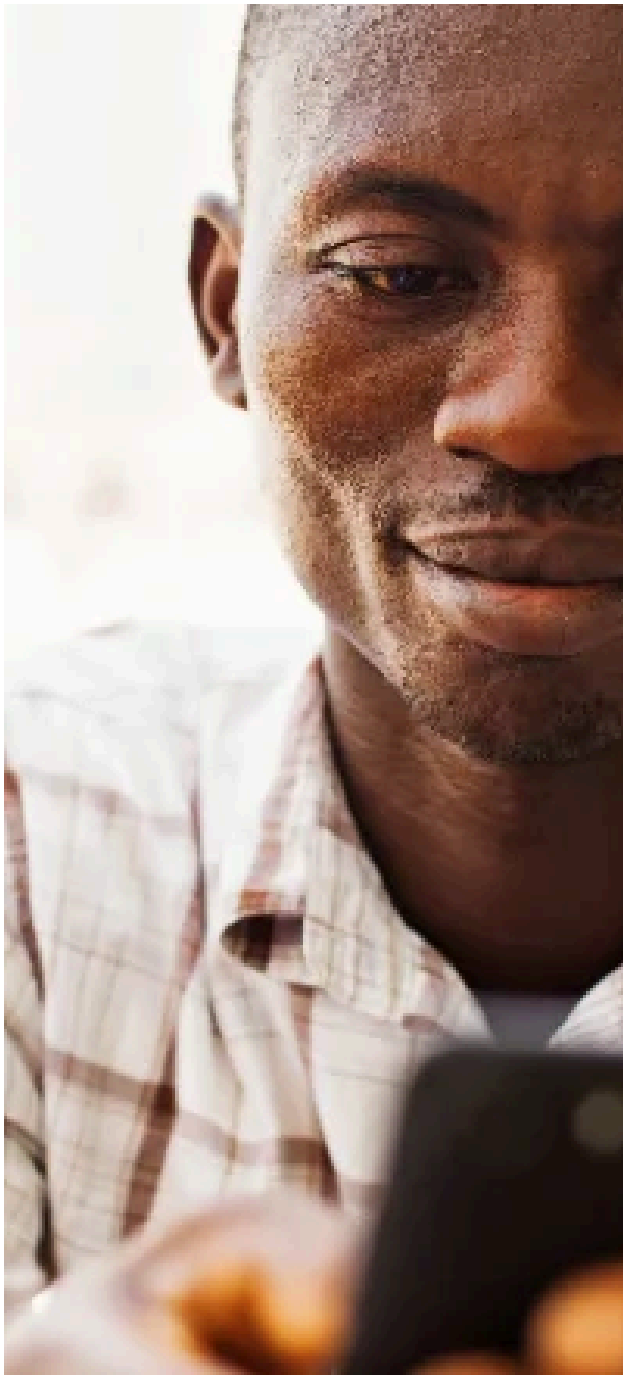
Introduction	Dr Simon Bransfield-Garth, Akili AI	3
Foreword	Vincent Kimosop, KNCCI	4
At a glance	Key figures across both events	5

PART 1 · SUMMARY REPORT

Summary	6
Four messages for business leaders	7
What we found	8
Survey highlights	9
Two perspectives, one conclusion	10
AI governance: the first 90 days and beyond	11

PART 2 · THE FULL REPORT

1	About the research	14
2	AI adoption: widespread, informal and moving into core work	15
3	Visibility: the shadow AI problem	17
4	Governance and accountability	18
5	Data protection, vendors and incidents	20
6	Capability, investment and confidence	22
7	Two vantage points: practitioners and business leaders	24
8	What this means for leaders	26
9	The regulatory outlook	27
10	Where to look for guidance and support	30
	About the partners	32



ABOUT THIS REPORT

Presented by Akili AI and the Kenya National Chamber of Commerce and Industry (KNCCI). Based on 214 completed responses to the Akili Snapshot AI readiness assessment, collected in the week of 16 October 2026.

Part 1 is written for business leaders and can be read in ten minutes. Part 2 gives the full results and the regulatory outlook. Status of regulation as at 19 October 2026.

INTRODUCTION

AI arrived in Kenyan business before anyone decided to let it in.



AKILI AI

AI is everywhere in most Kenyan business, whether it's creating a report, summarising a meeting or simply in a web search. In some Kenyan companies AI is creating tremendous value but in the vast majority, the combination of informal use and poor linkage to business objectives means it often creates as much work as it saves. Managers report being inundated by long reports full of "AI slop", every job vacancy has dozens of automated applications and employees are being asked to use tools they don't fully understand.

Most AI has not entered the workplace intentionally, rather it has been added silently to the tools people routinely use on their computers or even their phones. The marketing team using it to create promotions, the operations team reducing workload or the HR team scanning CVs. But with it comes responsibility – for privacy, data management, decision-making, cybersecurity and customer interaction, often the approach is informal at best, exposing organisations to significant reputational, business and even regulatory risk. This is why AI ownership cannot simply sit in the IT department – it is a whole company responsibility, ultimately owned by senior management or the Board, whether the organisation is a multinational or a 5-person marketing agency.

Presented by Akili AI and the Kenya National Chamber of Commerce and Industry (KNCCI), this report provides a snapshot of the state of AI readiness in Kenya. It highlights some of the gaps between AI adoption and governance of the AI already in use, and the risks that result from that usage. It also points at some of the immediate actions businesses can take to take better advantage of AI while managing their exposure.

Dr Simon Bransfield-Garth
Founder and CEO, Akili AI

FOREWORD

Kenya is ideally placed to take advantage of AI.



The Kenya National Chamber of Commerce and Industry is delighted to participate in this ground breaking report on the state of AI readiness in Kenya. Since late 2022, AI has become one of the most important topics of conversation in Kenyan business and, whatever we may think about AI, we cannot ignore it. Like the Internet, it is a profound change to the way we work and live our lives.

Kenya, and KNCCI members are ideally placed to take advantage of AI. Our highly digital economy, excellent levels of awareness and unifying technologies such as M-PESA mean that member companies have the opportunity to build AI into their businesses to accelerate growth, capture new markets and create new products and services for customers.

Building a vibrant and prosperous business community in Kenya requires us all to embrace the massive potential for our members and society at large, while managing the risks. We need to manage it well for the benefit of all. At KNCCI, we are delighted to be playing a role in that transition.

Vincent Kimosop
Head of Policy, Research and Advocacy, KNCCI



AT A GLANCE · KEY FIGURES ACROSS BOTH EVENTS

Almost everyone uses AI. Almost no one governs it.

99%

use AI in some form, from free tools to features built into everyday software

4 in 10

use it in core processes or customer-facing decisions such as credit, claims or service

3 in 4

have no AI policy in force

4 in 5

have no named owner for AI risk

1 in 6

has a documented impact assessment covering AI, although the Data Protection Act already applies

3 in 4

plan to expand their use of AI in the next year, yet 7 in 10 could not detect or respond to an AI incident

Source: AI readiness survey, October 2026; 214 completed responses across both events.
Figures combine both events; results for each event are in Part 2.

aigovernancekenya.com

AI is running ahead of the rules.

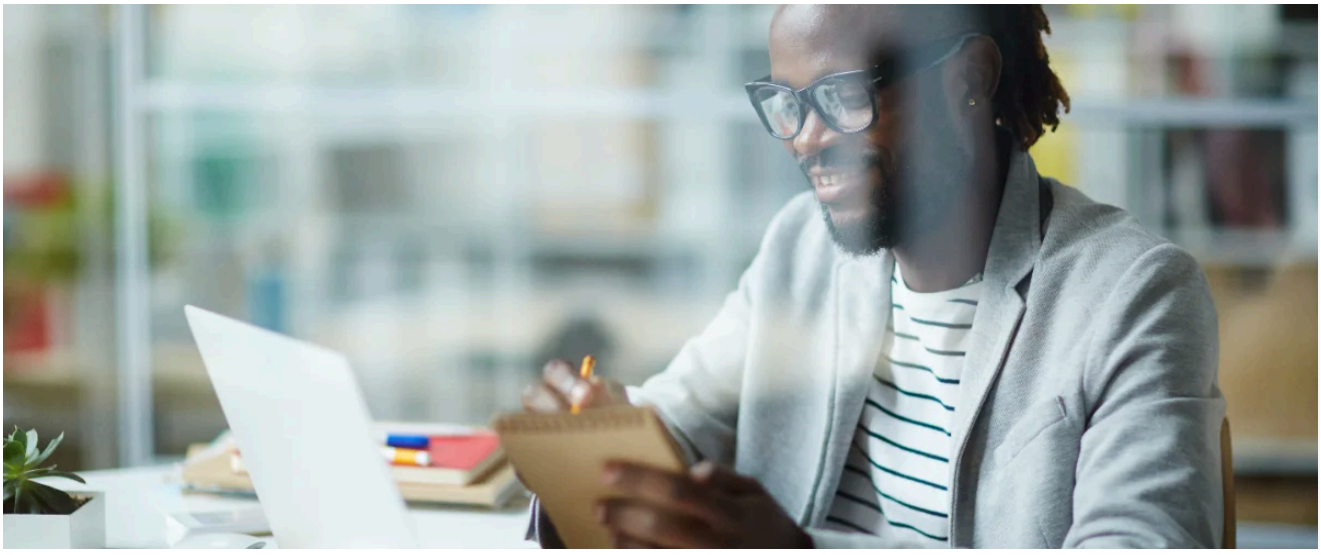
Findings from 214 respondents at the KNCCI Africa Commerce & Industry Summit and the Safaricom Cybersecurity Summit, and what business leaders should do in the next 90 days.

Summary

In the week of 16 October 2026, 214 people completed an AI readiness survey at two events. At the Safaricom Cybersecurity Summit, most were IT and security professionals from larger organisations. At the Africa Commerce & Industry Summit (ACIS), hosted by the Kenya National Chamber of Commerce and Industry (KNCCI), most were owners and executives of smaller businesses. The two groups see AI from very different positions. Their answers point to the same conclusion.

AI is already part of everyday work in almost every organisation surveyed, but governance has not kept pace. Most organisations cannot say what AI their people use, who is accountable for it, or where their data goes. This is not a problem that can wait for an AI law. Kenya's Data Protection Act 2019 already applies to personal data used in AI tools, and the data protection regulator is preparing AI-specific guidance.

The fixes are basic and achievable at any size. The following pages summarise the conclusions and set out a four-step plan for the first 90 days (measure, manage, train and evidence) and how to build on it. Part 2 gives the full results and the regulatory outlook for readers who want the detail.



Four messages for business leaders.

01

AI is already everywhere, including where firms think it isn't.

Almost every organisation uses AI, and four in ten use it in core processes or customer-facing decisions. Much of it arrives through staff using free public tools, or through features switched on in everyday software. Nearly one in five business leaders at ACIS gave answers suggesting they use no AI, yet all of them reported staff using public AI tools.

02

Training and governance are lagging.

Around two-thirds cannot see or control the AI their people use. Three-quarters have no AI policy in force, and most have no named executive accountable for AI risk.

03

Governance is not optional.

Existing data protection law already reaches how AI uses and stores data. The regulator is enforcing it and drafting AI-specific guidance, and wider AI regulation is on its way. Yet only one in six organisations has a documented impact assessment covering AI.

04

Action is more urgent than most realise.

Most organisations plan to expand their use of AI in the next year. Seven in ten would either miss an AI-related incident or have no plan to respond, and few could show a regulator how they manage AI today.

Where two figures are given, the first is for the Cybersecurity Summit and the second for ACIS.

1. AI is already everywhere

Nearly every organisation reports some AI use (99% of Cybersecurity Summit respondents and 100% of ACIS respondents). Four in ten already use it in core business processes or customer-facing decisions such as credit, claims or customer service. Much of it arrives informally, through staff using free public tools or AI features switched on in software bought for other reasons. Some leaders do not count this as AI at all: 18% of ACIS respondents answered at least one question as if they use no AI, yet all of them reported staff using public AI tools.

2. Training and governance is lacking

Around two-thirds either know staff use unapproved AI tools or have no way of telling (69% and 61%), and most hold no inventory of the AI in use. About three-quarters have no AI policy in force (76% and 77%). In larger organisations, AI risk usually sits informally with IT; at ACIS, 47% said nobody is accountable for it. Six in ten or more give staff no written rules on what data they may enter into AI tools, and around seven in ten would either not detect an AI-related incident or have no plan to respond.

3. Governance is not optional

Kenya has no AI-specific law yet: on 15 October 2026 the Senate ICT Committee recommended that the Artificial Intelligence Bill 2026 should not proceed until a national AI policy is in place. That does not leave AI unregulated. The Data Protection Act 2019 already governs personal data entered into AI tools, transfers of data to offshore AI services, and decisions made by automated means. The Office of the Data Protection Commissioner (ODPC) is actively enforcing the Act and has fined several organisations KES 5 million. Its draft Guidance Note on AI (July 2026) would require impact assessments for common uses such as credit scoring, recruitment screening, fraud detection and staff monitoring.

Few organisations are ready for this. Only 15% of Cybersecurity Summit and 23% of ACIS respondents have a documented impact assessment covering AI. Around two-thirds do not know where their AI vendors store their data, or whether it is used to train the vendor's models.

4. Action is more urgent than most realise

Optimism is high: 94% of Cybersecurity Summit and 98% of ACIS respondents see AI as more opportunity than risk. Most plan to expand their use in the next 12 months (80% and 65%), many without a policy in force. Security, privacy and regulation are already the most-cited barrier to getting value from AI, but in most organisations that concern has not yet turned into controls. Because seven in ten would either miss an AI incident or have no plan to respond, the first sign of a problem may be a customer complaint or a letter from the regulator.

AI use is running ahead of the rules.

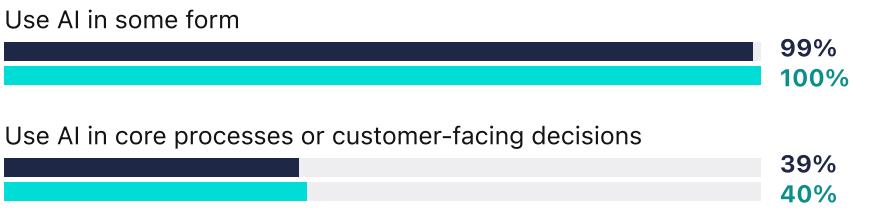
Survey of 214 respondents in the week of 16 October 2026.

Cybersecurity Summit: IT and security professionals (n=157)

ACIS: business owners and executives (n=57)

1 AI is already everywhere

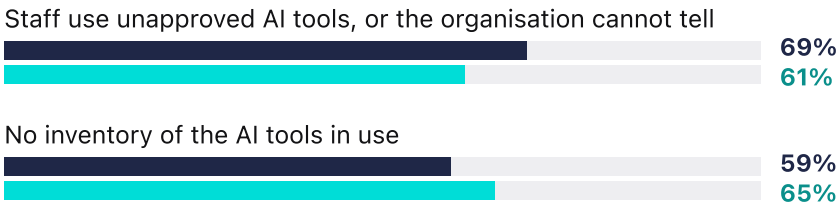
Use is near-universal. Much of it arrives through staff and everyday software, not through strategy.



18% of business leaders gave answers suggesting they use no AI, yet all of them reported staff using public AI tools.

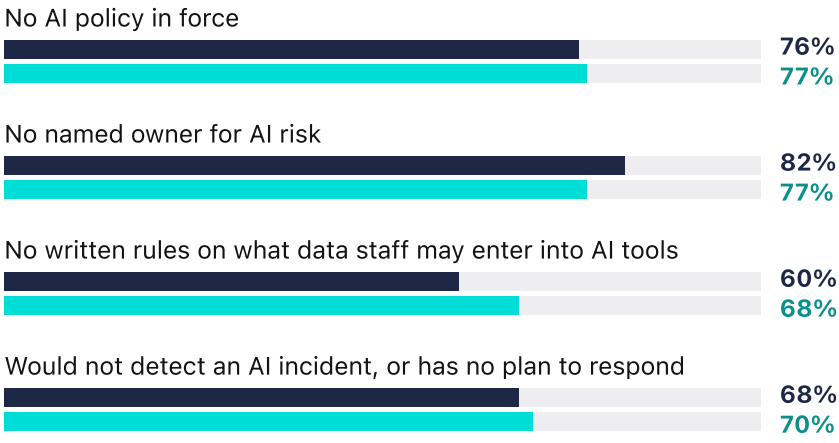
2 Much of it is shadow AI

Most organisations cannot see the AI their people are using.



3 Training and governance is lacking

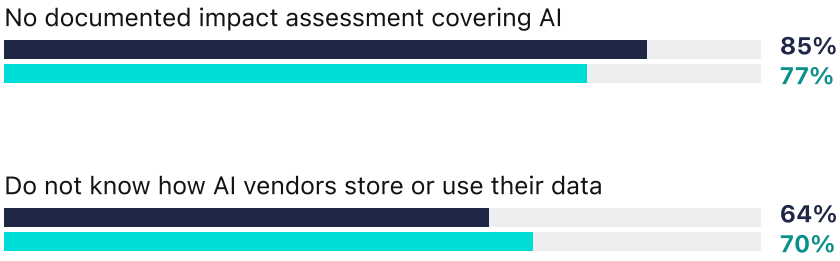
Basic controls are missing in most organisations, large and small.



Only 11% and 14% have all four basics: a policy in force, a named owner, written data rules and a data protection review.

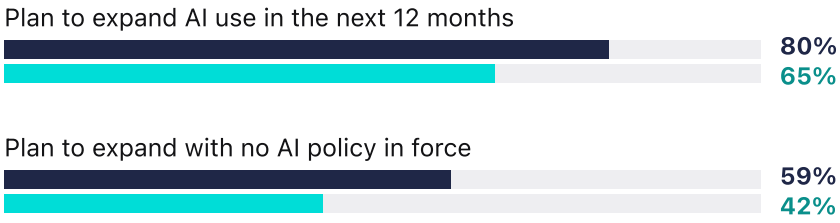
4 **The law already applies**

The Data Protection Act 2019 covers personal data used in AI tools today. AI-specific guidance from the regulator and a national AI policy are close.



5 **Act now, before growth outruns control**

Confidence is high and use is growing, mostly without a policy in force.



Source: AI readiness survey, October 2026. Cybersecurity Summit n=157 and ACIS n=57 completed responses. Percentages of completed responses.

PART 1 · TWO PERSPECTIVES, ONE CONCLUSION

SHARE OF RESPONDENTS	CYBERSECURITY	ACIS
Know staff use unapproved AI tools	43%	21%
Nobody is accountable for AI risk	12%	47%
AI risk is left informally with IT	70%	30%
Senior leaders do not understand AI, or rely entirely on others	40%	58%
Would trust AI output unchecked for most or all work	32%	44%

Practitioners see more of the problem; business leaders are more likely to be exposed without knowing it. At ACIS, the 58% on leaders' understanding is largely a self-assessment. Size is not destiny, though. At ACIS, 35% route all AI use through an approval process and 33% formally review new AI tools, showing that good governance is achievable in small firms.

AI governance: the first 90 days and beyond.

Good AI governance does not start with a long policy document. It starts with knowing what is happening, putting simple controls around it, making sure people can use AI safely, and keeping the records that show it. The first 90 days build that foundation; the following 6 to 12 months turn it into a full AI governance programme.

STEP 1

Measure

Find out what AI is in use, including unapproved tools and AI features switched on in licensed software. Record what data each use touches, where it goes and how much risk it carries.

Leads: CIO or head of IT

STEP 2

Manage

Name an executive owner who reports to the board. Issue a short policy with approved tools and data rules, assess high-risk uses against the Data Protection Act, and add AI to the incident plan.

Leads: CEO, AI owner, DPO

STEP 3

Train

Give every user practical safe-use training, and brief the board and leadership on AI risks and their duties.

Leads: HR and AI owner

STEP 4

Evidence

Keep the inventory, policy, impact assessments, training records and a log of significant AI-supported decisions. Give the board a first AI governance report.

Leads: AI owner and DPO

For smaller businesses

The same four steps apply in a lighter form: a one-page list of the AI tools in use, the owner as the accountable person, a one-page rule sheet for staff, a short team briefing, and a folder that keeps the records. Check the settings in free AI tools that control whether your data is used for training.

Where to get help

Your data protection officer and the ODPC's guidance notes are the first port of call. aigovernancekenya.com offers practical guides for Kenyan organisations, regulatory updates and a free five-minute readiness check.

What to watch

- The ODPC's final Guidance Note on AI, which would make impact assessments mandatory for common AI uses.
- The Central Bank of Kenya's draft Guidance Note on AI in the banking sector, in consultation until 7 November 2026.
- Cabinet approval of the National AI and Emerging Technologies Policy 2026, which proposes a central AI regulator and an AI safety institute.
- A redrafted AI Bill, following the Senate ICT Committee's recommendation of 15 October 2026.
- For firms serving EU customers: most EU AI Act obligations for high-risk AI systems apply from 2 December 2027.



About the survey

Responses are self-reported and the samples self-selected, so the results show patterns rather than national estimates. Percentages are of completed responses and are rounded; figures in the at-a-glance graphic combine both events. A full report with detailed results is provided in Part 2.

PART 2

The full report.

Survey results, implications and the regulatory outlook

Short of time? Part 1 has everything you need to act.

This part is for readers who want the evidence behind the conclusions: the detailed results, what they mean for different roles, the regulatory outlook and where to find further guidance.

About the research.

This report draws on 214 completed responses from two very different audiences. That allows a rare comparison between the people who run Kenyan organisations' technology and the people who run the organisations themselves.

Respondents completed the same 25-question AI readiness survey in the week of 16 October 2026: at the Safaricom Cybersecurity Summit, and at the Africa Commerce & Industry Summit (ACIS) hosted by the Kenya National Chamber of Commerce and Industry (KNCCI). The questions cover three areas: adoption and maturity, governance and data protection, and capability and confidence. Of 316 surveys started, 214 were completed; the 102 that were not completed contained no answers and are excluded.

Exhibit 1. Two contrasting groups of respondents

	Cybersecurity Summit	ACIS
Main roles	IT and technology 72%; cybersecurity 15%; executives 6%	Owners, executives and board members 72%; operations, HR and finance 18%
Organisation size	500+ staff 30%; 100–499 staff 45%; under 100 staff 25%	1–5 staff 58%; 6–24 staff 23%; 25+ staff 19%
Largest sectors	Banking and financial services 13%; government 11%; telecoms and technology 11%; manufacturing 10%	Agriculture 28%; retail and consumer 14%; professional services 12%
Organisation type	Private sector 78%; public sector and parastatals 15%; NGOs 7%	Private sector 82%; NGOs 12%; public sector and parastatals 5%

Source: AI readiness survey, October 2026.

How to read the figures

- Percentages are shares of completed responses and are rounded, so answers may not sum to exactly 100%.
- Where two figures are given, the first is for the Cybersecurity Summit and the second for ACIS. Combined figures cover all 214 respondents, of whom 73% were at the Cybersecurity Summit.
- The ACIS base is small: each respondent is worth almost 2 percentage points, so differences of a few points should not be over-read. Sub-groups of fewer than 50 respondents are flagged.
- Both samples are self-selected and self-reported. People attending a cybersecurity event are likely to be more aware of AI risk than average, so if anything the results may flatter the wider business population. The results show patterns, not national estimates.

The survey highlights in Part 1 summarise the main results for each event. Sections 2 to 7 take each theme in turn.

AI adoption: widespread, informal and moving into core work.

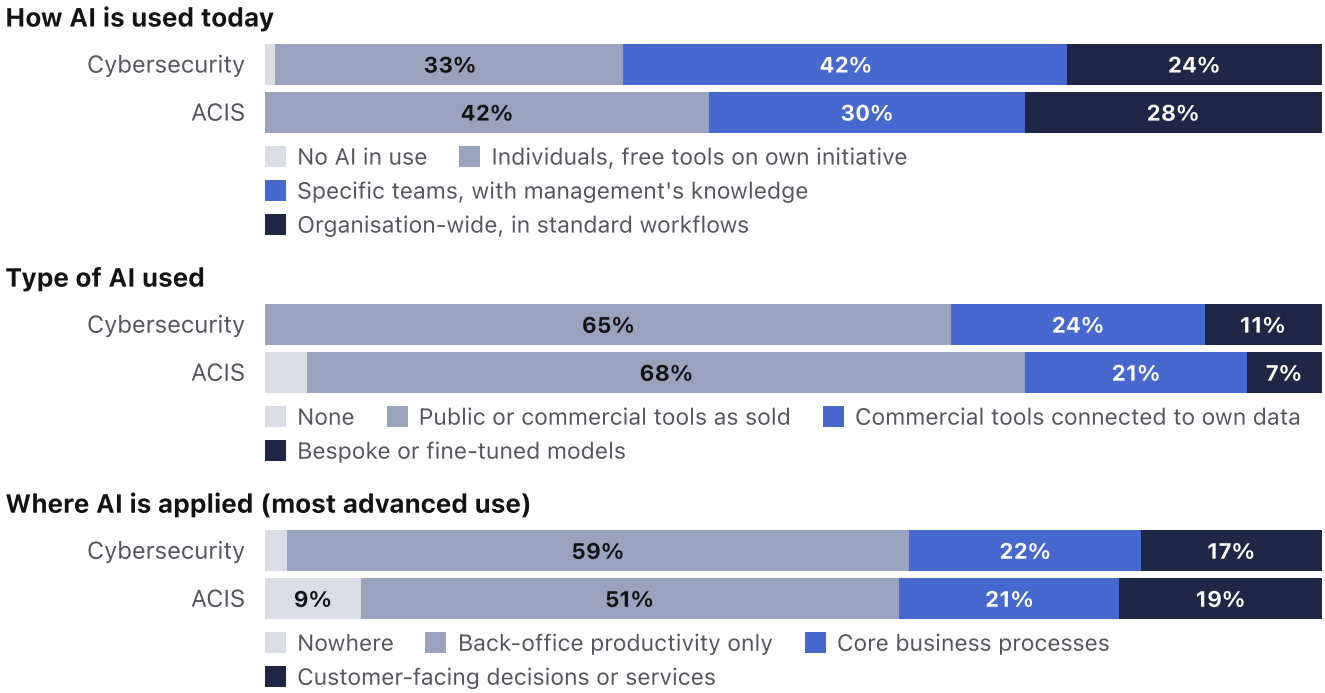
AI has arrived in Kenyan organisations through people rather than strategy, and it is already moving beyond the back office.

Almost every organisation reports some use of AI (99% and 100%). But it is rarely planned: a third of Cybersecurity Summit respondents and 42% at ACIS say AI use is driven by individuals using free public tools on their own initiative. Only around a quarter (24% and 28%) describe AI as part of standard workflows across the organisation.

Most use relies on off-the-shelf tools such as ChatGPT, Copilot or Gemini (65% and 68%). A significant minority has gone further, connecting commercial tools to their own documents, data or systems (24% and 21%) or using bespoke or fine-tuned models (11% and 7%). These deeper integrations raise the stakes, because the organisation's own data is now flowing into the AI.

AI is also moving into work that matters. While most use is still in back-office productivity (59% and 51%), around four in ten organisations already apply it in core business processes such as finance, HR or security monitoring (22% and 21%), or in customer-facing decisions and services such as credit, claims and chatbots (17% and 19%).

Exhibit 3. AI use is widespread, informal and moving into core work
Share of respondents by answer



Source: AI readiness survey, October 2026. Cybersecurity Summit n=157; ACIS n=57 completed responses.

AI that organisations do not count

Some AI is invisible to the people answering. 18% of ACIS respondents and 4% at the Cybersecurity Summit answered at least one question as if their organisation uses no AI, yet reported staff using public AI tools. AI also arrives inside software bought for other reasons: 84% and 65% say AI features are already active in tools such as Microsoft 365, Google Workspace, ERP or security software.

Adoption is accelerating. 82% and 84% say AI use has grown over the past year, and 11% at the Cybersecurity Summit say it has grown faster than they can manage.

WHAT THIS MEANS

AI exposure is not limited to organisations that have deliberately deployed AI. Any organisation whose staff use public tools, or whose software suppliers have switched on AI features, already has AI risk. An honest picture has to start from what people actually use, not from the list of approved projects.



Visibility: the shadow AI problem.

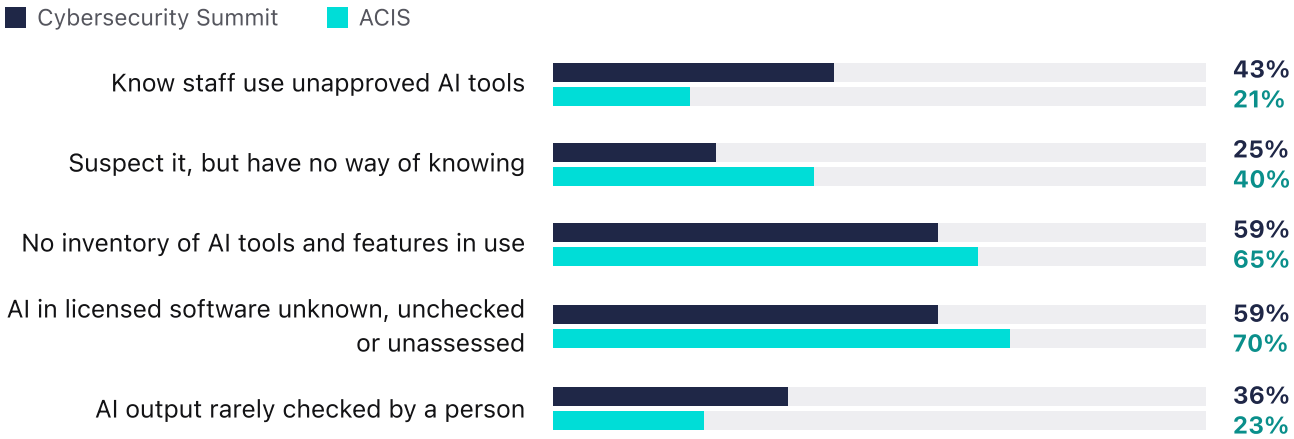
Most organisations cannot see the AI their people use, and that makes every other control harder.

Around two-thirds either know staff use AI tools that have not been approved (43% and 21%) or suspect it but have no way of knowing (25% and 40%). Only 20% at the Cybersecurity Summit and 35% at ACIS say all AI use goes through an approval route they monitor.

The gaps compound. Most organisations hold no inventory of the AI tools and features in use (59% and 65%), and most have not assessed the AI built into software they already license (59% and 70% either do not know, have not checked or have not assessed it). Once AI output is produced, checking is uneven: 36% and 23% say it is rarely reviewed by a person before use.

Exhibit 4. Most organisations cannot see the AI their people use

Share of respondents



Source: AI readiness survey, October 2026. Cybersecurity Summit n=157; ACIS n=57 completed responses.

The two groups see the problem differently. Practitioners are twice as likely to know shadow AI is happening (43% against 21%). Business leaders are more likely to have no way of knowing (40% against 25%), which suggests the problem is larger in smaller organisations than their owners realise.

WHAT THIS MEANS

Visibility comes first, because policies, training and assessments cannot cover what nobody knows is there.

Shadow AI is best read as a signal of demand rather than misconduct. Staff are using AI because it helps them. Experience elsewhere suggests that outright bans rarely work; providing approved, secure alternatives with clear rules does.

Governance and accountability.

The basic building blocks of AI governance are missing in most organisations, and size offers no protection.

Three-quarters have no AI policy in force (76% and 77%): many are drafting one (45% and 18%), and 31% and 60% have none at all. Only 7% at the Cybersecurity Summit and 16% at ACIS have a policy that has been issued, trained on and reviewed in the last year.

Accountability is the weakest point. In the larger organisations represented at the Cybersecurity Summit, AI risk usually sits informally with IT (70%). At ACIS, 47% say nobody specific is accountable. Only around one in five organisations (18% and 23%) have a named executive or risk owner, and fewer still (8% and 11%) one who reports to the board.

Controls on new tools and data are similarly thin. 42% and 61% have no meaningful gate before a new AI tool is adopted. 60% and 68% give staff no written rules on what data they may enter into AI tools, and only around one in five back written rules with technical controls that block restricted data. 60% and 54% keep no record of significant decisions made or supported by AI.

Bigger is not better governed

Among Cybersecurity Summit respondents, organisations with 500 or more staff were no better governed than smaller ones, and on accountability they were worse: 15% have an AI policy in force and only 6% a named owner for AI risk, compared with 26% and 21% in organisations of 100 to 499 staff. (Sub-groups of 47 and 70 respondents.) Complexity, many business units and informal IT ownership appear to hold larger organisations back.

Controls cluster together

Organisations with one control in place are much more likely to have others (Exhibit 6). Those with an AI policy in force are more than twice as likely to hold an inventory of AI in use, and far less likely to be unable to tell whether staff use unapproved tools. Organisations with a named owner for AI risk are three times as likely to say their incident process covers AI.

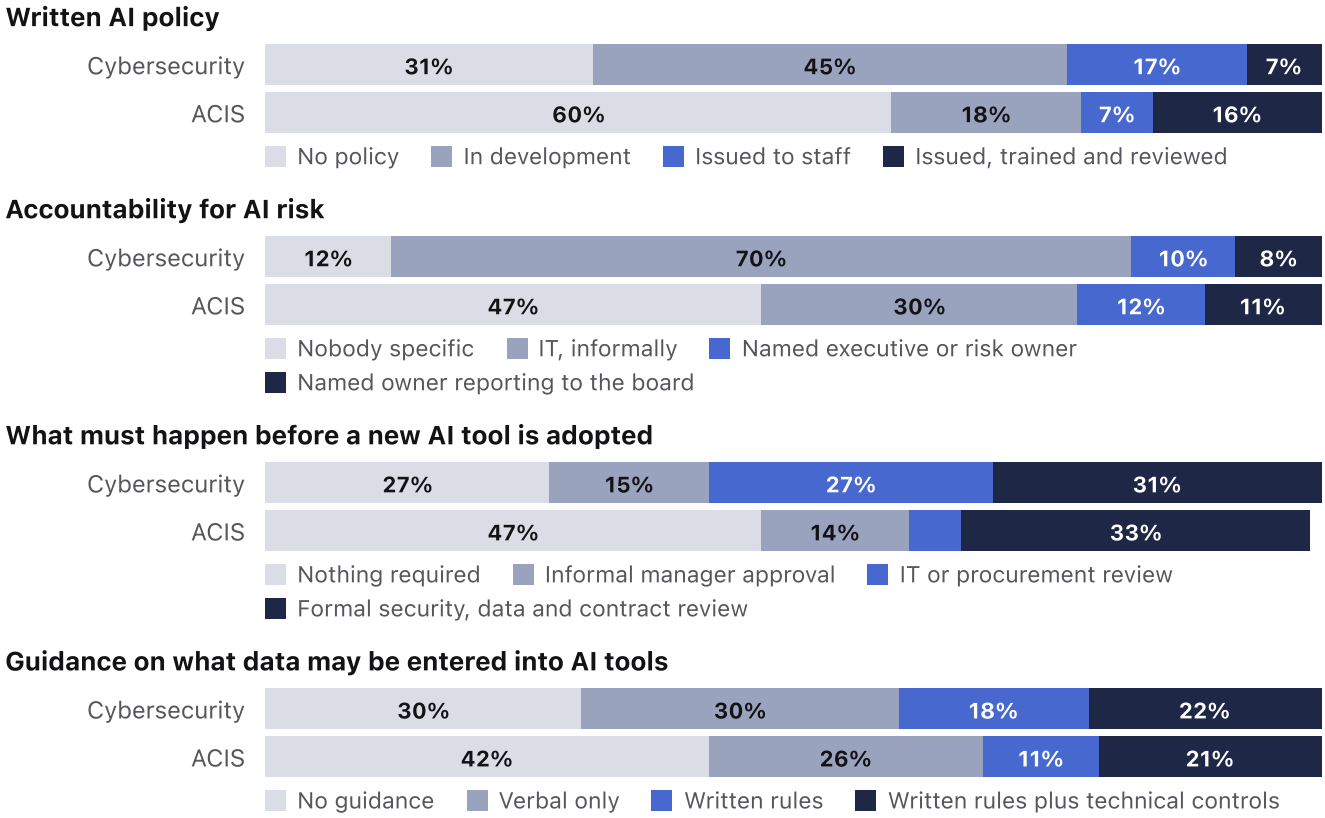
WHAT THIS MEANS

Governance does not have to start everywhere at once. A named, senior owner is the single most useful starting point, because the other controls tend to follow.

Only 12% of all respondents have the four basics in place: a policy in force, a named owner, written data rules and a review of data protection obligations. The gap between the leaders and the rest is wide, but the leaders show it can be closed.

Exhibit 5. The building blocks of governance are mostly missing

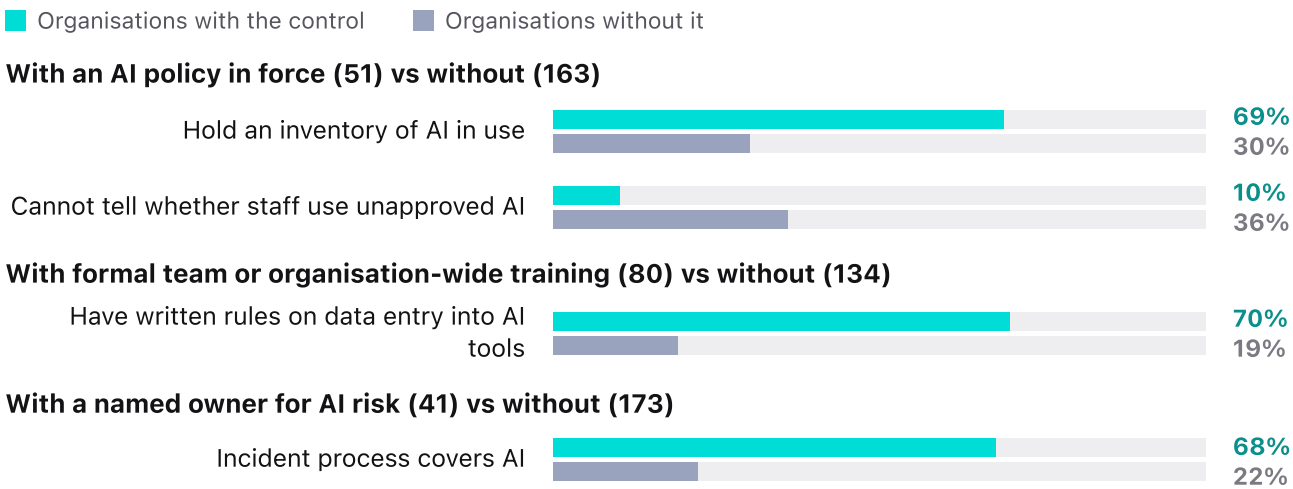
Share of respondents by answer, from least to most mature



Source: AI readiness survey, October 2026. Cybersecurity Summit n=157; ACIS n=57 completed responses.

Exhibit 6. Controls reinforce each other: one in place makes the next more likely

Both events combined (214 respondents); share of each group



Associations, not proof of cause: organisations that take governance seriously tend to do several things at once. Source: AI readiness survey, October 2026; both events combined, 214 completed responses.

Data protection, vendors and incidents.

The legal exposure is real today, and few organisations can evidence that they have addressed it.

Kenya's Data Protection Act 2019 already applies whenever personal data is entered into, processed by or produced by an AI tool. It requires a lawful basis for processing, transparency with individuals, impact assessments for high-risk processing, safeguards for transfers outside Kenya, written contracts with processors, and protection for individuals against significant decisions made solely by automated means.

Awareness of these obligations is uneven, and evidence is scarce. At the Cybersecurity Summit, 61% say they have reviewed their obligations in relation to AI, but most of them (46% of all respondents) only believe they comply; just 15% have a documented impact assessment covering AI. At ACIS, 54% have not considered their obligations at all, although 23% have a documented assessment.

Vendor risk is largely unexamined. 64% and 70% either do not know where data entered into their AI tools is stored and whether it trains the vendor's models, or simply assume it is handled properly. Only 10% and 14% have confirmed this for all tools in contract terms.

Incident readiness is low. 27% and 42% say they would probably not detect an AI-related incident such as a data leak, wrong decision or fraud; a further 41% and 28% might detect it but have no defined response. Only 8% and 21% have tested an AI incident scenario.

The highest-risk uses are the least documented

The 84 organisations already using AI in core processes or customer-facing decisions face the greatest exposure, because these are the uses for which regulators expect impact assessments and human oversight. Yet 67% of them have no AI policy in force, 77% have no documented impact assessment covering AI, and 44% keep no record of significant AI-supported decisions.

WHAT THIS MEANS

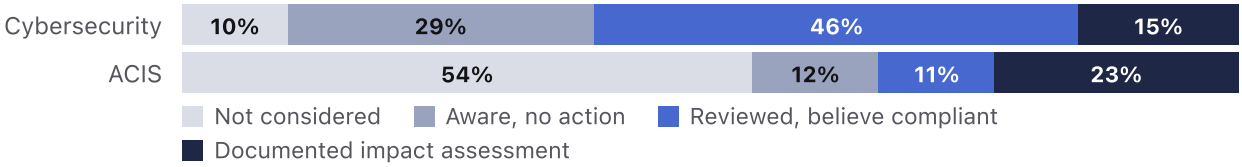
For regulators, documentation is the currency of compliance. An organisation that 'believes it complies' but cannot produce an impact assessment, a processor contract or a decision record is poorly placed if a customer complains or the regulator asks.

Generative AI tools used through personal accounts are a particular gap: data can leave the organisation, and possibly the country, with no contract in place at all.

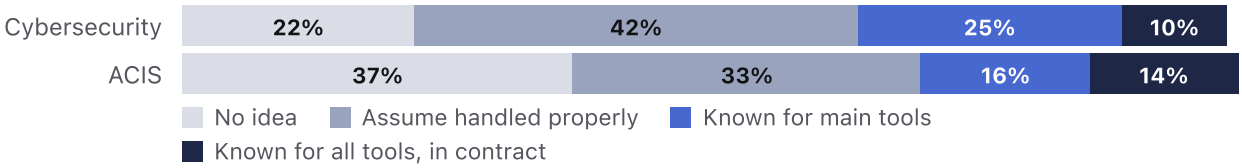
Exhibit 7. Data protection, vendor and incident readiness

Share of respondents by answer, from least to most mature

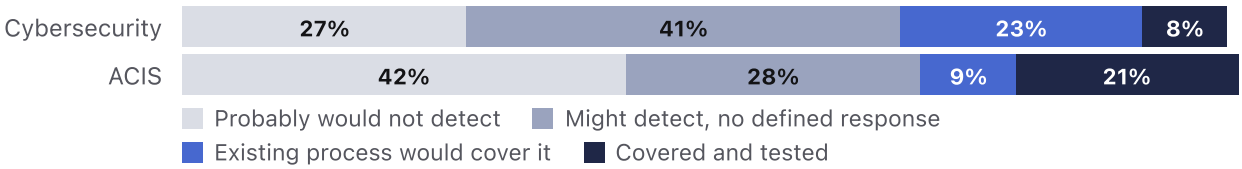
Data Protection Act 2019 obligations in relation to AI



Knowledge of where AI vendors store data and whether they train on it



Ability to detect and respond to an AI-related incident



Source: AI readiness survey, October 2026. Cybersecurity Summit n=157; ACIS n=57 completed responses.



Capability, investment and confidence.

Enthusiasm for AI is running well ahead of the skills, leadership understanding and budgets needed to use it safely.

Training is mostly informal: 63% in both groups provide no structured training, relying instead on self-learning by interested staff or nothing at all. Only 19% and 23% provide organisation-wide training that covers safe use and data handling.

Leadership understanding is a concern, particularly among smaller businesses. 40% at the Cybersecurity Summit say senior leaders do not understand AI or rely entirely on others to make decisions about it. At ACIS the figure is 58%, and because most ACIS respondents were themselves owners or executives, this is largely a self-assessment.

In-house capability and funding are thin. Over half of Cybersecurity Summit organisations (52%) have only general IT staff with no AI-specific skills; at ACIS, 35% depend entirely on vendors. Two-thirds or more have no defined AI budget (68% and 75%), and only 10% and 19% budget for security, governance and training rather than tools alone.

High confidence, unaddressed concern

Sentiment is overwhelmingly positive: 94% and 98% see AI as more opportunity than risk, and most plan to expand their use in the next year (80% and 65%). Around a third or more would trust AI output without checking it for most or all work (32% and 44%).

Security, privacy and regulation are the most-cited barrier to getting more value from AI (45% and 39%), ahead of skills (27% and 35%), data readiness (15% and 25%) and cost (13% and 2%). Yet among respondents who named security, privacy or regulation as their main barrier, 76% have no AI policy in force and 61% have no written rules on data use. The concern is real; the response has not yet followed.

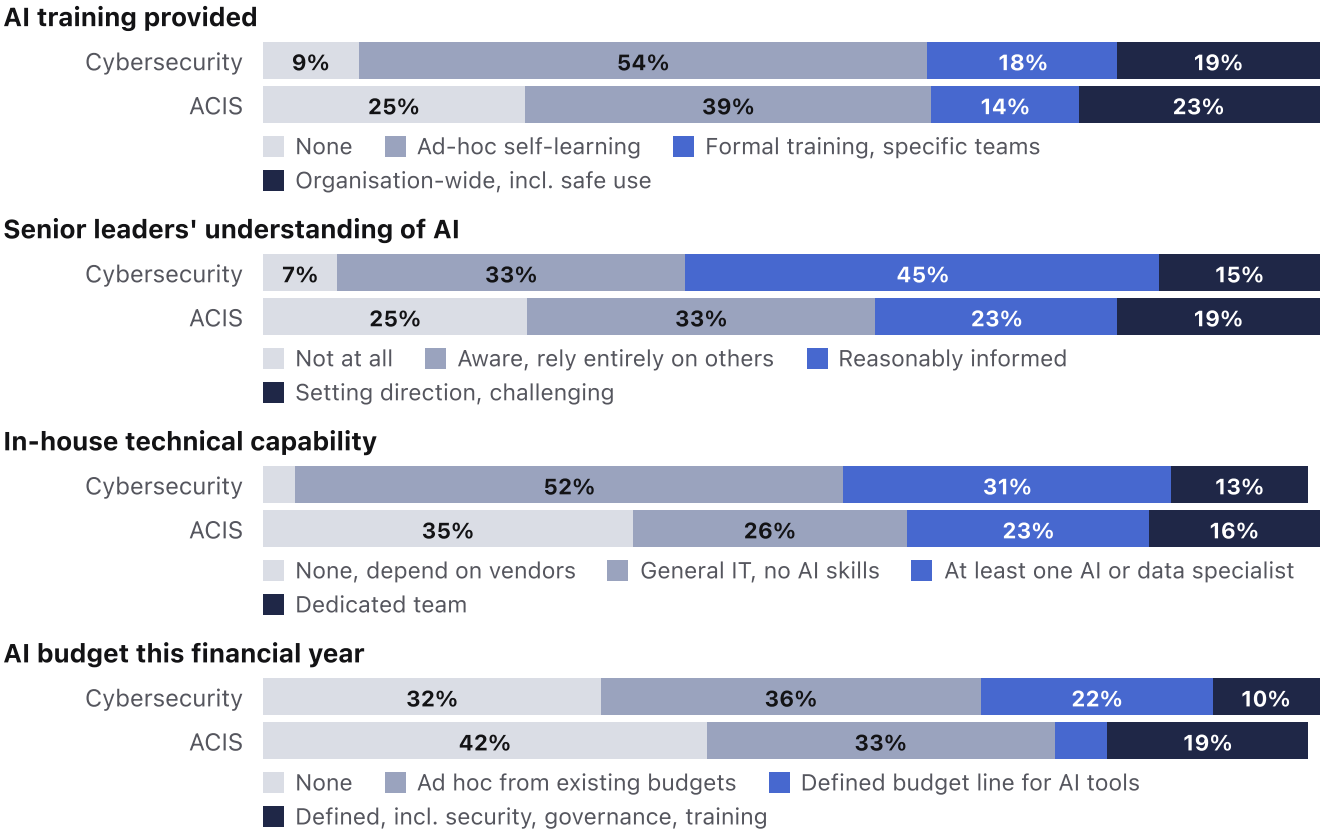
WHAT THIS MEANS

The capability gap is as much about leadership as technology. Boards and executives who do not understand AI cannot set risk appetite for it or challenge proposals.

Training pays off directly: organisations that provide formal training are more than three times as likely to have written rules on data use (70% against 19%).

Exhibit 8. Skills, leadership understanding and budgets lag behind use

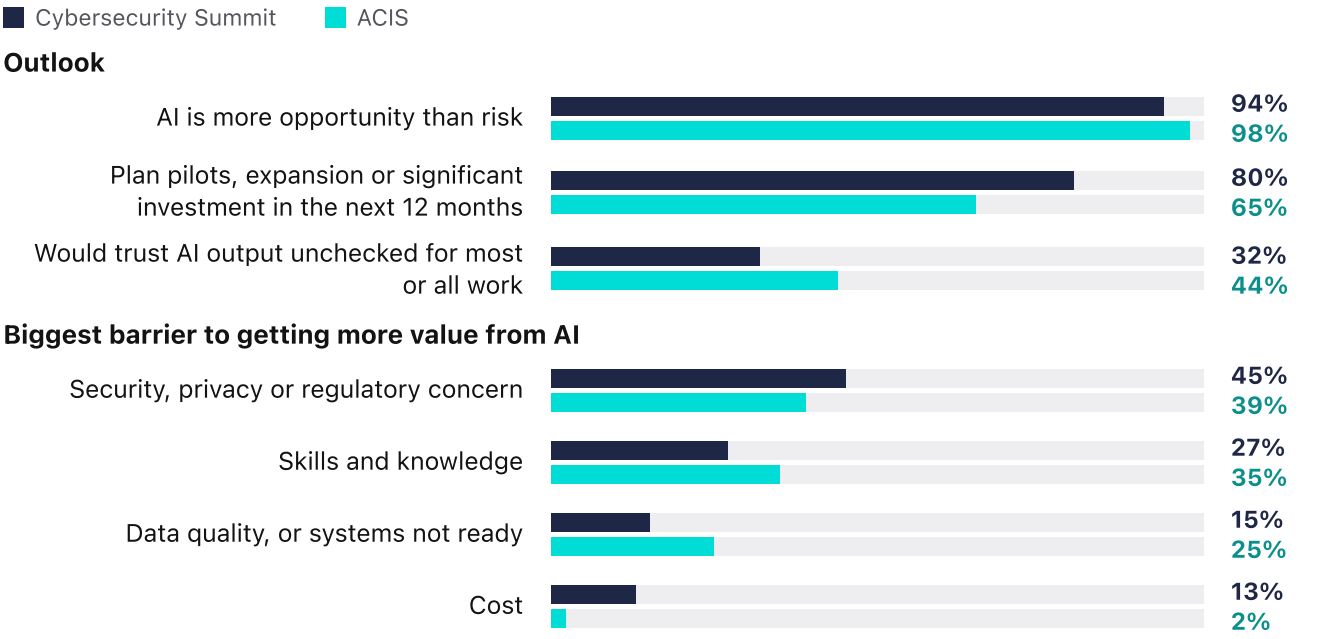
Share of respondents by answer, from least to most mature



Source: AI readiness survey, October 2026. Cybersecurity Summit n=157; ACIS n=57 completed responses.

Exhibit 9. Confidence is high; the main brake is security and regulation

Share of respondents

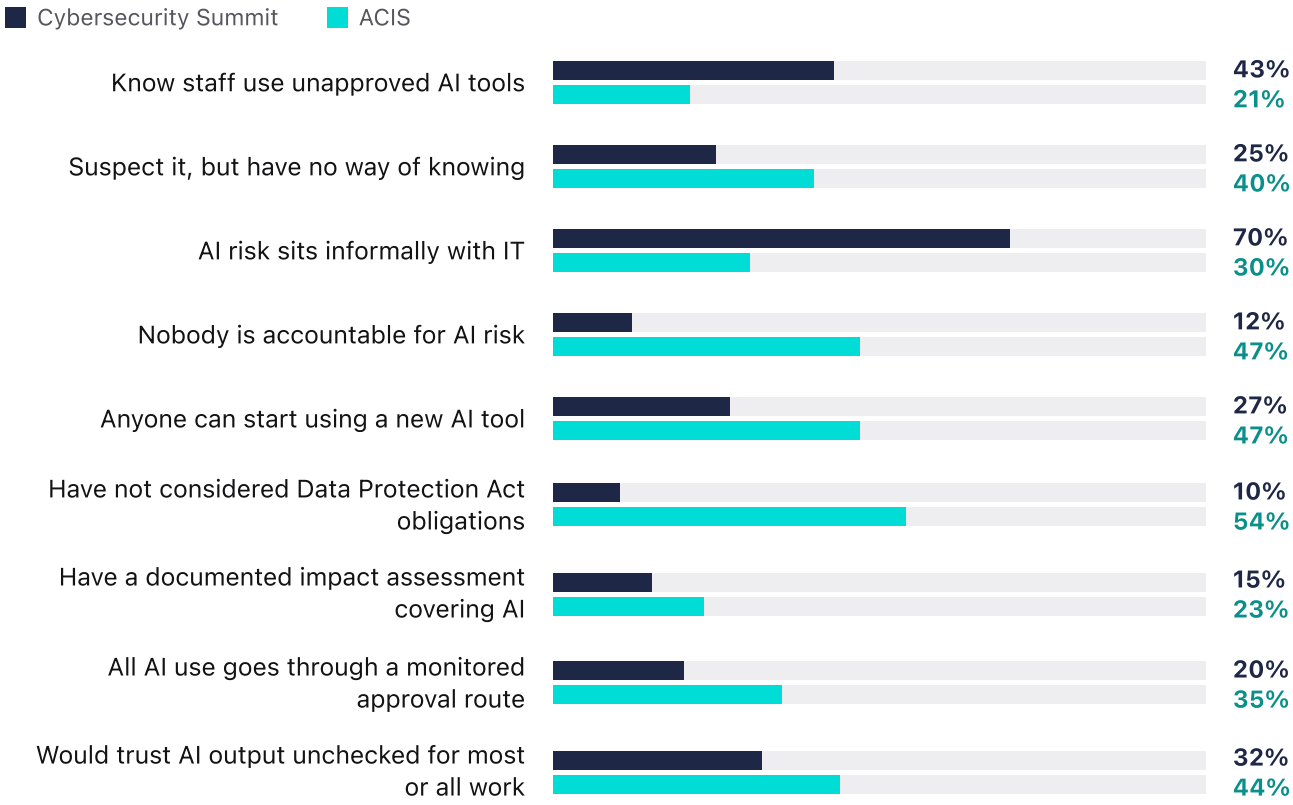


Source: AI readiness survey, October 2026. Cybersecurity Summit n=157; ACIS n=57 completed responses.

Two vantage points: practitioners and business leaders.

IT and security practitioners and business leaders describe the same gap from different sides, and neither group is well governed.

Exhibit 10. Where the two groups differ most
Share of respondents



Source: AI readiness survey, October 2026. Cybersecurity Summit n=157; ACIS n=57 completed responses.

Four observations

Practitioners see more; leaders control less.

The technology function is more aware of shadow AI, but in larger organisations it often carries AI risk informally, without the authority or board visibility to manage it.

Confidence on paper is not evidence.

At the Cybersecurity Summit, 46% believe they comply with data protection law in relation to AI without having documented it. ACIS respondents are less likely to have reviewed their obligations but slightly more likely to have documented them.

Small businesses are polarised.

A capable minority at ACIS routes all AI use through approval (35%), formally reviews new tools (33%) and has tested an AI incident scenario (21%), while a large group has no gate at all (47%) and has not considered its legal obligations (54%). Good governance does not depend on size.

Trust and checking do not line up.

Business leaders are more willing to trust unchecked AI output (44% against 32%), yet also more likely to say all output is reviewed (39% against 24%). Stated practice and actual behaviour may differ, which is a reason to measure rather than assume.



What this means for leaders.

The findings carry different implications depending on where you sit, but the common thread is the same: the exposure exists now, and it can be managed with basic, affordable steps, set out in Part 1.

Boards and chief executives	AI risk is an enterprise risk, not an IT matter. In most organisations surveyed, nobody at board level is accountable for it. Three questions are worth asking now: what AI do we use, including in our suppliers' software; who owns the risk; and could we show the regulator our impact assessments and decision records if asked?
Chief information and security officers	Shadow AI and AI features switched on by suppliers have widened the attack surface and created new routes for data to leave the organisation. AI tools should be treated as third parties that process data: brought into vendor risk management, covered by data-loss controls, and included in incident response plans and exercises. The technology function also needs the mandate to act, which means a named executive owner above it.
Data protection officers, legal and compliance teams	Expect a backlog of impact assessments, particularly for uses the ODPC's AI guidance treats as high risk: credit scoring, recruitment screening, fraud detection, staff monitoring and customer-facing decisions. Priorities include processor contracts with AI vendors, the lawful basis for cross-border transfers, human review of significant automated decisions, and records that show how decisions were made.
Owners of smaller businesses	The same law applies regardless of size, but the response can be proportionate: knowing which AI tools staff use, one accountable person, a one-page rule sheet, checking the settings that control whether free tools train on your data, and keeping simple records. The ACIS results show many small firms already do this well.
Policymakers, regulators and industry bodies	The results suggest that awareness and practical guidance are the immediate need, especially for smaller firms, where 54% have not considered their data protection obligations in relation to AI. Sector guidance from existing regulators, such as the ODPC's AI guidance note and the Central Bank's draft guidance for banks, is likely to have more immediate effect than new legislation; templates aimed at small businesses would help.

The regulatory outlook.

Kenya does not need an AI Act for AI to be regulated. Data protection, cybercrime and sector rules already apply, and AI-specific expectations are arriving through existing regulators faster than through new legislation.

Already in force

- **Data Protection Act 2019 and its 2021 regulations.** These cover the lawful basis for processing, transparency, data subject rights including protection from solely automated decisions, impact assessments for high-risk processing, conditions for transfers outside Kenya, registration of data controllers and processors, and breach notification. The ODPC is actively enforcing: it reports more than 2,000 complaints since 2022 and has imposed fines of KES 5 million on several organisations.
- **Computer Misuse and Cybercrimes Act.** Amended in 2025, with the amendments in force since November 2025 (some provisions have since been challenged in court). It is relevant to AI-enabled fraud, impersonation and misuse of computer systems.
- **Sector rules.** Requirements on technology risk, cybersecurity, outsourcing and fair treatment of customers from the Central Bank of Kenya, Capital Markets Authority, Insurance Regulatory Authority and Communications Authority apply to AI in the same way as to other technology. Consumer protection and employment law apply where AI influences decisions about customers or staff.

In development

- **ODPC Guidance Note on Artificial Intelligence.** Published in July 2026 and consulted on in August, it sets a minimum standard for organisations that develop, deploy or procure AI using personal data. It makes impact assessments mandatory for uses including credit scoring, health diagnostics, facial recognition, fraud detection, recruitment screening, student assessment, content recommendation and employee monitoring, and requires a documented review for others such as generative chatbots. It also addresses offshore AI processing, vendor contracts and human review of automated decisions.
- **Central Bank of Kenya draft Guidance Note on AI in the banking sector.** In consultation from September until 7 November 2026, for institutions licensed under the Banking Act. Other financial regulators may follow.
- **National AI and Emerging Technologies Policy 2026.** Consultation closed on 4 August 2026. The draft proposes a central regulator, a Kenya AI Safety Institute and a public AI registry, and is expected to go to Cabinet next.
- **Artificial Intelligence Bill 2026.** A Senate bill proposing an AI Commissioner, registration of AI systems and penalties of up to KES 5 million. On 15 October 2026 the Senate ICT Committee recommended that it should not proceed, citing breadth, penalties and overlap with existing regulators, and asked for the national policy to be finalised first.
- **Standards.** The Kenya Bureau of Standards' code of practice for AI applications (KS 3007) sets out non-binding expectations on risk management, data governance, testing, monitoring and incident response.

International

- **EU AI Act.** It applies to organisations whose AI systems are placed on the EU market or whose outputs are used in the EU, so it reaches Kenyan exporters and service providers with EU customers. Prohibited practices have applied since February 2025 and duties for general-purpose AI models since August 2025. Following the 2026 Digital Omnibus, most obligations for stand-alone high-risk systems apply from 2 December 2027 and for high-risk AI in regulated products from 2 August 2028.
- **African Union Continental AI Strategy (2024).** A framework for member states that emphasises responsible, inclusive adoption of AI.
- **Voluntary standards.** ISO/IEC 42001 (AI management systems) is becoming the reference point for customers and partners who ask how suppliers manage AI.

Sources: ODPC; Central Bank of Kenya; Ministry of Information, Communications and the Digital Economy; Parliament of Kenya; Official Journal of the EU; African Union. Status as at 19 October 2026.

Exhibit 11. The regulatory picture is filling in quickly

Selected Kenyan and international milestones for organisations using AI in Kenya



What to expect

Regulation through existing institutions.

The Senate committee's preference for building on existing regulators, and the ODPC and CBK guidance already issued, point to sector-led, risk-based oversight rather than a single new AI regulator in the near term.

Documentation as the test.

Impact assessments, processor contracts, decision records and evidence of human oversight are what regulators will ask to see.

Customers and partners before regulators.

Banks, multinationals and international buyers are increasingly asking suppliers how they govern AI, often before any regulator does.

A moving target.

Guidance will be finalised and extended over the next 12 to 24 months. Organisations that build basic governance now will adapt more easily than those that wait for final rules.



Where to look for guidance and support.

Much of the help organisations need already exists. The list below is a starting point, grouped by type.

TYPE	RESOURCE	WHAT IT OFFERS
Practical guidance	**aigovernancekenya.com**	Practical guides on AI governance for Kenyan organisations, updates on regulation and policy.
Inside your organisation	**Data protection officer**	First point of contact on personal data in AI: impact assessments, vendor terms and data subject rights. If you have no DPO, name someone responsible.
	Internal audit and risk committee	Independent assurance that AI risks are identified and controls work; a route for AI risk to reach the board.
Regulators and public bodies	**Office of the Data Protection Commissioner (ODPC)**	Registration, complaints, guidance notes including on impact assessments, and the Guidance Note on AI.
	Central Bank of Kenya (CBK)	For banks and payment service providers: guidance on cybersecurity and outsourcing, and the draft guidance on AI.
	Capital Markets Authority (CMA)	Rules for capital markets firms, and a regulatory sandbox for testing innovative products.
	Insurance Regulatory Authority (IRA)	Rules for insurers on outsourcing, conduct and treatment of customers.
	Communications Authority and National KE-CIRT/CC	National coordination of cyber incidents, threat advisories and incident reporting.
	Ministry of Information, Communications and the Digital Economy	The National AI Strategy and consultations on AI policy.
	Kenya Bureau of Standards (KEBS)	KS 3007, the Kenya Standard code of practice for AI applications.

TYPE	RESOURCE	WHAT IT OFFERS
Frameworks and standards	**ISO/IEC 42001**	International standard for AI management systems; certifiable, and increasingly requested by customers.
	NIST AI Risk Management Framework	Free, practical framework for identifying and managing AI risks, with a profile for generative AI.
	OECD AI Principles	Widely adopted principles for trustworthy AI, useful for framing a policy.
Research and policy networks	**KICTANet**	Multi-stakeholder ICT policy network; analysis of Kenya's AI policy and legislation.
	CIPIT, Strathmore University	Research on AI, data protection and technology policy in Kenya and Africa.
Business networks	**KNCCI**	Business support, events and advocacy through national and county chambers.
	Kenya Private Sector Alliance (KEPSA)	Private sector engagement with government on policy, including digital and AI.

Inclusion in this list does not imply endorsement. Links checked in November 2026.



About Akili AI

Akili AI is a Kenyan-led artificial intelligence company that helps organisations adopt AI productively and responsibly, from AI readiness assessments and governance to building and running AI applications. Its Founder and CEO, Dr Simon Bransfield-Garth, and Co-founder and CTO, Shikoli Makatiani, together bring decades of experience in technology and AI.

The Akili Snapshot is a 25-question AI readiness assessment covering adoption and maturity, governance and data protection, and capability and confidence. A free five-minute readiness check is available at aigovernancekenya.com.

About KNCCI

The Kenya National Chamber of Commerce and Industry is the voice of Kenyan business. For 60 years it has advanced trade and industry, providing business support, events and advocacy through national and county chambers, and generating evidence to inform policy and support its members.

KNCCI hosts the Africa Commerce & Industry Summit (ACIS), where the business-leader responses in this report were collected.

Acknowledgements

Our thanks to the delegates at the Africa Commerce & Industry Summit and the Safaricom Cybersecurity Summit who completed the survey, and to the teams at KNCCI and Akili AI who made the research possible.

Citation

Akili AI and KNCCI (2026). *AI Readiness in Kenya 2026: AI is running ahead of the rules.* Nairobi, November 2026. © 2026 Akili AI and the Kenya National Chamber of Commerce and Industry. Responses are self-reported; results show patterns rather than national estimates.

START NOW

You cannot get value from, or manage, what you cannot see.

Most leaders guess they have 15 to 20 applications using AI. The answer is usually 50 to 100. Closing the gap is not difficult. It just has to be done.

MEASURE
Find the AI in use, the data it touches and the risks

MANAGE
Name an owner, set the rules, approve safe tools

TRAIN
Equip every user and leader to use AI safely

EVIDENCE
Keep the records that prove you are in control

Practical guidance, regulatory updates and a free five-minute readiness check

aigovernancekenya.com