Busara Groundwork
THOUGHT PIECE

Schomerus, Mareike[1]

NO 36
November 2026

Busara's AI strategy 2026 – 28: Skilled, value-driven, and equitable AI interaction for stronger research, advisory, and practice

When the currency of an organization is knowledge, how this currency is minted matters.



Busara

Key words:
AI strategy
AI values
AI use cases
social science research
AI transparency

Author affiliations:
[1] Busara

Busara seeks a reflective approach to transparency on knowledge production. We use the Contributor Role Taxonomy (CRediT) (https://credit.niso.org) to make visible the many inputs that producing research requires beyond the named authors. We combine CRediT with the sIfA tool to also transparently reflect on how humans and AI interacted in the production of an output.

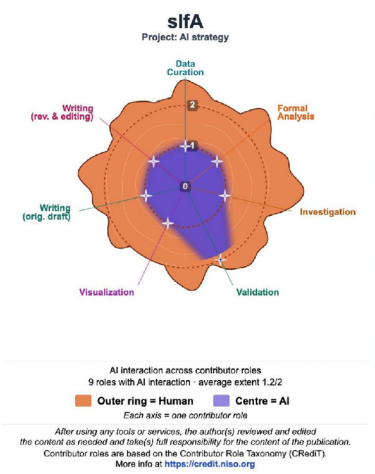
Conflicts of interest:
There are no conflicts of interest to declare for this study.

Copyedited by:
Hitha Em

Designed by:
Lynette Gow

Acknowledgements
This strategy benefited tremendously from the generous input and time commitment of Busara staff who attended the consultations or shared their views in other ways. Thanks go in particular to the Busara Futures Initiative co-pilot Engy Saleh as well as to Konrad Gulla and Dan Nommensen for their guidance and teaching.

Statement of Intellectual Fellowship and Accountability (sIfA)



Statement of Intellectual Fellowship and Accountability (sIfA) — AI interaction across contributor roles

CRediT Role	Performed by	Extent of AI interaction (0 none, 1 some, 2 extensive)	How was AI used? Audit trail?	Tool / Service
Data Curation	Mareike Schomerus *Sole*	1	Note taking in online calls ; background tech research	Claude Sonnet 5
Formal Analysis	Mareike Schomerus *Sole*	1	Summary assessment of all transcripts from consultations (minus persona development); return on investment calculations	Claude Sonnet 5
Investigation	Mareike Schomerus *Sole*	1	see audit trail	Claude Sonnet 5
Validation	Mareike Schomerus *Sole*	2	Uploaded the segment requirements that TÜV Austria requires for a certified AI strategy, asked AI to check my drafted text against the requirements with segment per segment validation, see audit trail	Claude Opus 5
Visualization	Mareike Schomerus *Sole*	1	Swimming Lanes graphs	Claude Opus 5
Writing – original draft	Mareike Schomerus *Sole*	1	Minimal connective prose (mainly in the sentiment analysis); summary writing of exec summary from existing prose only	Claude Sonnet 5
Writing – review & editing	Mareike Schomerus *Sole*	1	see audit trail	Claude Sonnet 5

This sIfA was created using: Schomerus, M. (2026). sIfA Tool - Statement of Intellectual Fellowship and Accountability (Version 1.4) [Computer software]. Busara. https://doi.org/10.5281/zenodo.20285992

The sIfA tool can be used and downloaded at www.busara.global or at https://github.com/mareikeschomerus-ctrl/sIfA.

Released under the Apache License 2.0

Generated 2026-09-13 17:13

See Appendix: sIfA for a more detailed audit trail

Figure 1: sIfA for the AI strategy

Abbreviations and acronyms

AUDAS	Artificial intelligence (here used broadly to include generative and non-generative)
BD	Busara's project framework: Align, Understand, Design, Assess, Share
BizOps	Business Development
BR-UK	Business Operations. One of Busara's five core divisions. Behavioural Research UK
BRACE	Busara's Behavioral Research and Academic Engagements department
CGD	Center for Global Development
COM-B	Capability, Opportunity, Motivation = Behaviour; prominent behavior change model
COPE	Committee on Publication Ethics
COSTAR	Context, Objective, Style, Tone, Audience, Response; widely-taught prompt-engineering framework
CRediT	Contributor Roles Taxonomy; community-owned taxonomy for giving credit in research work
CREME	Former Busara department dealing with research ethics, now part of MRI
DAIR	Distributed AI Research Institute
DPDP	Digital Personal Data Protection Act (India, 2023)
DPIA	Data Protection Impact Assessment
FARM	Busara's Food Systems, Agriculture and Climate Resilience Management
FIVE	Futures, Impact, Voice and Engagement, one of Busara's five core divisions in 2026
FTE	Full-Time Equivalent
GDPR	General Data Protection Regulation (EU)
GenAI	Generative AI
IRB	Institutional Review Board

Abbreviations and acronyms

ISERC	Institutional Scientific and Ethics Review Committee
IT	Information Technology
KM	Knowledge Management
KPI	Key Performance Indicator
LGPD	Lei Geral de Proteção de Dados; Brazil's general data protection law
LLM	Large Language Model
MRI	Methods, Research and Innovation; part of the Research Excellence department at Busara
MVP	Minimum Viable Product
NACOSTI	National Commission for Science, Technology and Innovation (Kenya)
NIST AI RMF	US National Institute of Standards and Technology AI Risk Management Framework
PAIE	Project AI Evidence
RAG	Retrieval Augmented Generation
RAP	Research, Advisory and Practice; one of Busara's five core divisions
RFP	Request for Proposals
ROI	Return on Investment
RPA	Robotic Process Automation
SaaS	Software as a Service
SCP	Social Coherence Portfolio in Busara
sIfA	Statement of Intellectual Fellowship and Accountability; Busara's common good AI-disclosure tool for human/AI interaction
SLT	Senior Leadership Team
STM	International Association of Scientific, Technical and Medical Publishers
Tafakari	Busara annual report

Table of contents

Executive summary 8
Introduction: Why does Busara need an AI strategy? 18
AI vision and goals 25
Value proposition of AI 44
Industry-specific applications 45
Stakeholders and change management 52
Solutions and resources 67
Technologies and tools 82
Data strategy 88
Security and ethics 95
Revenue potential and cost efficiency 104
Implementation plan 107
References 114
Appendices 116
Appendix 1: PROTECT Code of Busara's SCP 116
Appendix 2: Busara's internal context 117
Appendix 3: Mandate 117
Appendix 4: sIfA 123

Tables

Table 1: Aligning org-strategy KPIs with sections in the AI strategy 21
Table 2: How do org KPIs map onto AI strategy KPIs? 40
Table 3: Broad comparison of how Busara's peers approach their AI strategy? 50
Table 4: Use case suitability check 72
Table 5: What do different technologies do for Busara's use cases? 83
Table 6: Retrieval/Knowledge-base platform options 86
Table 7: Qualitative scenarios for avoiding losses from budget overruns 105
Table 8: Annual value recovered through time saved on information searching 106
Table 9: Evidence-generating projects 107
Table 10: What processes exist and what is needed 109
Table 11: Activities plan for AI strategy development 118

Figures and Boxes

Figure 1: sIfA for the AI strategy 3

Box 1: The spirit of what Busara seeks in an AI strategy 56
Box 2: Busara was built to counter bias. Yet, AI bias is not a major topic 61
Box 3: Busara has a trust challenge 63
Box 4: Tension: Can we not just improve our win rate by submitting many more AI-assisted proposals? 71
Box 5: Tensions in strategy alignment and prioritization 100

Executive summary

This AI strategy sets out Busara's approach to how we use and interact with AI. It is rooted in the stance that AI offers both tremendous opportunity as well as an existential threat to social science research. To make the most of the opportunity and to mitigate the threat requires a clear AI vision and a values-based approach to our learning on what AI needs to mean for our work.

How to navigate this document

This document offers a comprehensive insight into Busara's thinking on how to approach AI. This approach includes a full commitment to transparency on challenges the organization faces, and which Busara believes are shared by others. Being open about these is, in Busara's view, essential to finding constructive ways to approach AI and to support future collaborations. However, we see that it is a risk to be transparent about the challenges organizations face. Yet, we do believe that collectively, research organizations like Busara need an open conversation about these.

The AI strategy follows ISO/IEC 17024 requirements for what a certified AI Manager needs to deliver as an organizational AI strategy. As such, this document might serve not only to articulate Busara's approach, but can hopefully inspire others to take a comprehensive and thought approach to navigating the challenge of AI.

Busara's AI vision

Busara's AI vision is to strengthen our partnerships and technical quality by optimizing internal processes and moving research questions on AI use center-stage to become an innovative, skilled and recognized learning leader on the role of AI for ethical research with Global Majority populations.

Our aim with this vision is to remain sustainably competitive by including all elements of Busara (such as the overall strategy, the organization's structure and processes, and its culture) in an iterative, agile, unfragmented, and learning-driven approach to AI. We achieve this by using AI to translate our values into goals on critical thinking, AI governance as well as improving processes and outputs. This requires us to build knowledge, skills and behaviors of how Busara staff interact with AI where it supports this vision, and maintain AI-free work where it does not.

Five bold steps in our approach to AI

These five bold steps are the north star that guides us through both high-level strategic decisions on how to use AI as well as daily interactions with AI. They are designed to allow people to ask of each other, their AI-related decisions and whether those are true to the five bold steps.

1 We do not automate the mind

We strengthen human thinking (analysis, attention, problem solving, open-mindedness) and decision making on substantive points (project management, research design, networking and business development) by avoiding decision fatigue through AI-supported processes that add little quality or innovation value (budget tracking, proposal logistics, opportunity scanning). This means we need to be bold in developing automated processes that we trust, and intentionally focus our mind.

2 We defend our guard rails

We will work towards agreed and explicit rules on how humans interact with AI in research, writing, proposals, and other processes. We continue our learning to allow these rules to evolve in ways that are relevant and appropriate to the ethos of our work. This means that we say no to efficiency if it increases speed without increasing quality or supporting innovation. We need to be bold in defending our guard rails when our competitors prioritize speed.

3 We build Busara's institutional brain

Every proposal, research project or networking opportunity benefits from having institutional knowledge readily accessible. This means we build technology to facilitate this access. We do this while boldly seeking to change the current culture and behaviors on knowledge management, which are informal and messy.

4 We refuse to let AI flatten our voice

As a champion of diverse styles of expression and epistemologies, we actively counter a flattened AI voice, even when it is less efficient or difficult to do. This means we prioritize our integrity over efficiency and mass appeal. We put in the work to find and develop our voice and we boldly defend that Busara outputs do not sound like everyone else's.

5 We write the rules on AI governance and transparency before someone else writes them for us

We are honest about the challenges we are facing when it comes to AI and about what procedural flaws we are seeking to address. We are vulnerable and radically honest in disclosing our AI use. We name our governance principles. We publish research on the role of AI in Global Majority social sciences. This is a bold step because it is emotionally challenging to admit when we have flaws and what help we sought to deal with them. But this bold step positions Busara as thought leader on AI and knowledge production. And it is the right thing to do.

↘ Six goals for our AI strategy

Goal 1: Build AI-literate, critical thinkers
Goal 2: Establish responsible AI governance
Goal 3: Identify and operationalize AI opportunities for 'one RAP (Research, Advisory and Practice) way' for systematic delivery
Goal 4: Utilize AI to improve delivery of diverse, high-quality, and operationally-relevant publications

Goal 5: Diagnose where AI can help fix flawed processes and data practices
Goal 6: Utilize AI to improve networking and business development

Ten priority AI use cases

Use case 1: Critical thinking coaching
Use case 2: AI-related project work inventory
Use case 3: Transparency and AI governance
Use case 4: Meeting notes and action-item tracking
Use case 5: Project management
Use case 6: Building a queryable institutional memory that enforces data protection and prioritizes AI on the input side
Use case 7: Tagging/classification support for the outputs map
Use case 8: AI-assisted public adoption alerts
Use case 9: First-pass literature screening and lit-review support
Use case 10: Proposal/BD support

Why does Busara need an AI strategy?

As a research organization by and for the Global Majority, Busara is facing a number of external and internal challenges:

- Funding pressure due to shrinking aid budgets and shifting funding priorities;
- Tough competition as peer organizations and funders are visibly moving on integrating AI into work or requests for proposals;
- Increasing requests for clarity on our stance towards AI (these include being pushed to use it more or the request to use it less);
- Demand to build an AI-skilled workforce;
- Already inequitable skills development, which is exacerbated by inequitable AI use, putting additional pressure on our need to continue to build the skills of researchers and business operations staff, including cutting-edge AI skills;

- A substantial gap and backlog on how we manage our knowledge and utilize it for our work going forward;
- Governance and risk exposure (including risk to our rigor) from ungoverned shadow AI, driven by job loss fears and general fear of missing out on AI developments;
- An unclear organizational culture on how to discuss and address uneven skill sets or AI use, creating debates that can be emotive (driven by a few) or incomplete (because of a lack of more general participation);
- Motivational fallout from feeling internationally sidelined due to external funding crunches and diminishing collaborations.

This AI strategy helps Busara meet these challenges while enhancing its value in the following ways:

- For Busara, low-threshold access tiered to different levels of AI skill supports equitable upskilling, streamlines and improves processes, opens access to knowledge, and carries set-piece work so innovation can happen where it matters. Because growth can be a consequence of impact, this approach could increase the evidence and learning asset rate and total revenue, contributing to org-key performance indicators (KPIs).
- For our partners, the AI strategy creates a Busara that meets the moment: anticipating change with innovation, and resolving long-standing issues of governance, quality, skills, values and processes while delivering more reliably and sustainably on its core work of human knowledge creation, contributing to org-KPIs.
- The broader societal benefit lies in sharing our learning and governance of AI for ethical research with the Global Majority as a common good, to improve global practices and outcomes for all.

Which Busara values and expectations underpin this strategy?

In addition to Busara's known organizational values, the following expectations and values came to the fore during consultations for the strategies.

Expectations:

- No forced AI usage: the option to remain AI-free;
- Not simply policing of AI use, but constructive approaches driven by ethos and values;
- Thoughtful AI policies that build existing skills and add new ones;
- Finding a way to counter high levels of distrust of colleagues and leadership.

Valued ways of working in Busara:

- Thinking: this value emphasizes the importance of challenging thinking, rigorous knowledge and the ability to deal with contrasting views.
- Different voices: this value emphasizes the need to produce diverse and contextually-rich work that represents not just Busara's multicultural team, but the many different places where the work is done. This value extends to championing different writing and speaking styles.
- Openness to challenge: Substantive disagreement, rooted in research, is welcome, but must be articulated rigorously, rather than expressed as a discomfort.
- Rigor and depth: Rigor and contextualized depth of the work are high aspirations for Busara and simultaneously two of the greatest pain points, with much work considered as lacking both.
- Learning: As a value, the organization champions giving people the opportunity to develop skills by doing the work, making mistakes, and improving.

What is explicitly not the goal of the AI strategy?

The AI strategy is not the final word on Busara's approach to AI. This is not a document that says: here is exactly how we will do it. Instead, the strategy offers a roadmap of necessary decisions on how to best fulfill the goals and how to design the use cases while staying true to the five bold steps.
Based on consultations with leadership and staff, it is not the goal of the strategy to:

- Turn Busara into an AI-first organization;
- Settle on long-term platform lock-in: staff is keen that we stay flexible and avoid vendor commitment that cannot be reversed should the politics of a provider strongly counter our values;
- Touch, reopen or reform existing ethics governance beyond adding clauses that speak specifically to AI developments that were not yet in scope when the ethics governance was articulated.

What happens if the AI strategy succeeds?

Success for the AI strategy is that within the implementation timeframe (which matches the 2026 - 2028 timeframe of the Busara 'Practice Partnerships for Impact' strategy), Busara is an AI-skilled research organization with improved processes and visibility.

Success means that the organization is known for thoughtful and transparent use of AI, a vibrant debate culture on AI, and publications that contribute to knowledge on the role of AI in research and particularly on how AI is impacting social science research with and by the Global Majority.

This also means that Busara is a role model to show what it looks like in practice if an organization leads in unpacking what the AI visions issued by governments or regional bodies mean for the future of skills building, values, and sovereignty.

Concretely, success means:

- Busara staff are skilled in working with AI (and consciously and collectively do not work with AI where this contradicts Busara's values, or does not add to innovation and quality, or degrades skills).
- Trust is restored where staff transparently disclose their AI use to each other, trust leadership to do the same, and trust leadership to have taken a rigorous, collaborative and thoughtful approach to making Busara a leading representative of responsible AI use in research with the Global Majority.
- AI use is contextually-appropriate and transparent. Both are hallmarks of Busara.
- Busara has an active discussion and learning culture on AI, with strong leadership on AI implementation across the organization.
- Busara's guardrails on AI usage are substantive and value driven. AI is used to make our work better, with broader knowledge and deeper debate, and not just faster.
- Busara is making a public contribution to questions on AI governance for knowledge production beyond championing internal capability: it is positioned and visible as a thought leader on aspects of knowledge production.
- Busara's knowledge management (KM) system is genuinely searchable and in active use.
- Staff have strong job-relevant skills and are developing AI skills. The latter are used where appropriate and always in combination with the organization's AI governance and risk assessment processes.
- Busara continues to speak with a diverse voice—or rather, with many voices.

Relevant legislation and frameworks

A number of legislative and strategic frameworks are relevant for this strategy:

- Kenya's 2026 AI Bill (not yet passed into law) covers risk classification, human rights impact assessment, cybersecurity requirements, ethics guidelines on bias, data privacy, environmental impact assessment. Kenya's draft also includes the stipulation that AI 'enhance rather than replace human capabilities', with mandatory human oversight for decisions affecting rights, safety, or wellbeing (p.32), which acts as a strong regulatory anchor for Busara's AI governance.
- Kenya's AI strategy emphasizes the need for Kenyan organizations to put effort into upskilling and reskilling towards home-grown solutions using local data and talent.
- The EU AI lists practices against which a behavioral research organization needs to actively screen: manipulation of behavior, exploitation of vulnerable persons, social scoring, and specific biometric identification uses (e.g., racial profiling), alongside emotion inference in the workplace or behavioral targeting that could be classified as seeking to manipulate behavior. Article 50 stipulates duties on transparency and how these apply to emotion recognition, biometrics and generated content.

Mandate and methods

This AI strategy was formally commissioned by Busara's CEO Francis Meyo with full authority for its development delegated to the author Mareike Schomerus, VP Futures, Impact, Voice and Engagement. To deliver the AI strategy first required external capacity building of the author, staff time of the co-lead of the Futures Initiative Engy Saleh (Senior Director BRACE), consultations space (org-wide meetings, consultation sessions, surveys), and executive team commitment and input as well as desk research, and peer feedback.

It is a living document, but will be signed off for completeness, feasibility and investment decisions by the CEO, after a final organizational feedback round on the full strategy.

Six phases of org-wide consultations (March through August 2026) helped capture a number of sentiments about AI. Early sessions surfaced AI use as a 'dirty secret in Busara; used but undisclosed. This disclosure gap prompted Busara Futures to design and ship the open-source sIfA tool, which uses existing community standards to declare human/AI interaction in knowledge production (www.sifatool.org). By April, staff named concrete use cases (coding, risk scoping, reviewing health results) and voiced hopes and fears for their own thinking. Through June, emotional responses ranged from AI enthusiasm to shame over labor exploitation in model training. By July, consultations shifted to co-designing solutions and some try outs on a platform. Separate executive team sessions raised a distinct concern: uneven adoption, not AI itself, is the primary organizational risk.

Introduction: Why does Busara need an AI strategy?

When the currency of an organization is knowledge, how this currency is minted matters.

The speed and force with which generative artificial intelligence (AI from here onwards) has arrived in our professional lives as researchers and policy thinkers has catapulted this insight to the foreground, putting unprecedented demands for transparency, clarity and intentionality on our minting process, if Busara is to maintain credibility, remain true to the ethos of social science research and heed the call of national and international AI legislation to deploy AI ethically, in ways that are appropriate to its context and that support growing a workforce of AI skilled professionals.

While some interpret this moment as one of additional and unwanted pressures to respond to a technological acceleration, our response is different. With this strategy, Busara seeks to acknowledge that while we have to find ways to strategically deal with accelerated technology, the demands this brings offer Busara an opportunity to comprehensively examine our processes and address a number of identified pain points within them. We do this while embracing transparency and clarity as the backbone of how we govern our knowledge production and our identity as a research organization with and for the Global Majority. This is a moment in which Busara also needs to clarify its values as a research organization and how it approaches AI, and use this stance to guide our internal processes and how we present ourselves externally.

Why now?

This strategy is needed now because as a research organization by and for the Global Majority, we are facing an existential challenge due to a number of internal and external compound drivers.

Externally, we are facing:

- Funding pressure due to shrinking aid budgets and shifting funding priorities;
- Tough competition as peer organizations and funders are visibly moving on integrating AI into their work or requests for proposals;
- Increasing requests for clarity on our stance towards AI (these include being pushed to use it more or the request to use it less);
- Demand to build an AI-skilled workforce.

Internally, the quality of our knowledge production work is threatened by:

- Already inequitable skills development, which is exacerbated by inequitable AI use, putting additional pressure on our need to continue to build the skills of researchers and business operations staff, including cutting-edge AI skills;
- A substantial gap and backlog on how we manage our knowledge and utilize it for our work going forward;
- Governance and risk exposure (including risk to our rigour) from ungoverned shadow AI, driven by job loss fears and general fear of missing out on AI developments;
- An unclear organizational culture on how to discuss and address uneven skill sets or AI use, creating debates that can be emotive (driven by a few) or incomplete (because of a lack of more general participation);
- Motivational fallout from feeling internationally sidelined due to external funding crunches and diminishing collaborations.

This is a good moment to implement a comprehensive AI strategy as Busara is just embarking on its 2026-28 strategy (org strategy from here on), which allows feeding AI strategy considerations into the broader strategic picture.

Who needs this AI strategy?

This strategy takes as its axiom that AI is pervasive across all professional profiles within Busara, and thus requires a comprehensive response. This means that all divisions—Executive Team; Research, Advisory and Practice (RAP) including Voice and Engagement (VE); Research Excellence (including Futures); Business Operations (BizOps)—need a strategic approach to AI. Within all divisions, this need stretches from the most junior roles to the CEO.

Busara's strategy and strategic goals

Busara's 2026 - 2028 strategy 'Practice Partnerships for Impact' highlights an ambitious growth path through long-term practice partnerships that 'keep behavior central to development outcomes by ensuring behavioral evidence translates into routines and systems that hold under real-world constraints.' This is clarified by the definition of 'practice' as 'Busara's disciplined work of helping institutions turn behavioral insight into operational routines that make the right change stick under real constraints such as turnover, politics, resource scarcity, and delivery pressure'.[1],[2]

Busara aims to achieve this strategy through two high-level goals:

1. ***'Fewer, larger, multi-year practice partnerships, so that at least 60% of revenue comes from renewable partnerships aligned to our three Practice areas [Health and well-being delivery systems; climate, environment and agrifood adoption systems; inclusive finance and economic resilience] by 2028'.***[3]

1 (Busara 2026)
2 (Busara 2026)
3 (Busara 2026)

2. *'Mainstreaming behavioral science through embedding behavioral and contextual insight into how real systems run, including decision rules, workflows, incentives, service design, accountability, and feedback loops'.*[4]

The organizational KPIs correspond to sections of the AI strategy in the ways outlined in Table 1 below. How organizational KPIs connect to the AI strategy is listed in.

Table 1. Aligning org-strategy KPIs with sections in the AI strategy

KPI	AI strategy corresponding section
KPI 1 — Total revenue	Value proposition; Technologies and tools; Data strategy
KPI 2 — Gross profit margin	Data strategy; Solutions and resources; Revenue potential and cost efficiency
KPI 3 — Practice partnership share	Stakeholders and change management; Data strategy; Industry-specific applications
KPI 4 — Scaled adoption rate	Implementation plan; Technologies and tools; Data strategy; Security and ethics
KPI 5 — Evidence and learning asset rate	Value proposition; Technologies and tools; Data strategy

4 (Busara 2026)

Busara's ambition is to be a research organization that delivers and promotes excellent, contextualized, ethical, and operationally-relevant behavioral research on and with the Global Majority. This is to inform better practices to support the ambitions, needs and human development of the people with whom the research was conducted. To deliver on this promise requires staff who are highly-skilled and supported by organizational processes that are efficient, transparent, functional, and reliable.

The AI strategy touches all aspects of Busara, informing decisions regarding budget allocation (for example on technology or training), platform procurement, personnel planning (through hiring policies that emphasize certain skills or reallocating staff focus to support the implementation of the AI strategy), and guardrails regarding what is considered acceptable AI use.

This AI strategy outlines a comprehensive approach that prioritizes the human role and agency with the principle of 'AI in the loop'. The concept of 'AI in the loop' is attributed to Stephen Roberts who argues 'that AI, rather than being the dominant factor in a loop of observation to action-taking, is actually relegated to being a support structure to help the human chain of command'.[5]

The aim is further to build a learning atmosphere that empowers staff in their thoughtful response to AI developments. This moment of technological acceleration can be a catalyst to address procedural shortcomings in Busara while strengthening our mission and the rigor of our work. A thoughtful AI approach has the potential to strengthen the quality of our work (through innovation and doubling down on essential research skills); allowing for partnerships to be negotiated transparently regarding AI use; and support a transformation to a systemic Busara, where organizational processes strengthen the core work of knowledge production for better practice.

5 (Bowman 2026)

What is explicitly not the goal of the AI strategy?

The AI strategy is not the final word on Busara's approach to AI. This is not a document that says: here is exactly how we will do it. Instead, the strategy offers a roadmap of necessary decisions on how to best fulfill the goals and how to design the use cases while staying true to the five bold steps.

Based on consultations with leadership and staff, it is not the goal of the strategy to:

- Turn Busara into an AI-first organization;
- Settle on long-term platform lock-in: staff is keen that we stay flexible and avoid vendor commitment that cannot be reversed should the politics of a provider strongly counter our values;
- Touch, reopen or reform existing ethics governance beyond adding clauses that speak specifically to AI developments that were not yet in scope when the ethics governance was articulated.

What does the AI strategy deliver?

The AI strategy delivers a comprehensive document that establishes where Busara currently is with regards to AI use, and where it wants to go, how it wants to interact with AI or not in specific processes, with what tools, to what value gain, and with what ethical and security considerations.

The planned outcome is for the strategy to set us up for implementation. However, its aim is also to galvanize conversations on values and ways of working and reforming processes. It seeks to strengthen a culture of knowledge sharing and learning to inform the strategic journey of Busara more generally, and specifically in the coming three years.

What happens if the AI strategy does not succeed?

An unsuccessful AI strategy has consequences: it likely will negatively impact Busara's reputation if shadow AI use continues with lack of clarity of how much of Busara's work was influenced by AI interaction. This will have a knock-on effect on discussions with funders and our ability to be taken seriously with our outputs. An unsuccessful strategy will also mean that neither leadership nor staff are held accountable for using AI responsibly, ethically, and in pursuit of stronger skills in their respective areas of work. This will happen if the AI strategy is not embedded in everyday processes, Busara's culture, ethics, and strategic ambitions.

> ***If the AI strategy does not succeed, Busara runs a risk of being seen as ignoring national, regional and international debates on AI visions. This could dampen Busara's ability to attract top talent.***
>
> ***More immediately, not having an AI strategy in place will continue current practices of shadow AI.***

What supports us in reaching our future?

We are supported by often highly-motivated staff. While there are often implementation challenges, we also have a climate in which learning is cherished. External pressures on funding and relevance bring to the foreground that we must meet this moment in a value-driven and learning-focused way to stay afloat.

AI vision and goals

Busara's AI vision statement

The AI strategy treats this moment of accelerating technology as an opportunity to audit where Busara must improve: our skills, processes, decision-making, technical quality of our work, thought leadership, data care, and our commitment to reflection and transparency. We assess first and use AI where it supports these improvements long-term and apply clear guardrails where it does not.

Our AI vision is to strengthen Busara's partnerships and technical quality by optimizing internal processes and moving research questions on AI use center-stage, to become an innovative, skilled, and recognized learning leader on the role of AI for ethical research with Global Majority populations.

Our aim with this vision is to remain sustainably competitive by including all elements of Busara (such as the overall strategy, the organization's structure and processes, and its culture) in an iterative, agile, unfragmented, and learning-driven approach to AI. We achieve this by using AI to translate our values into goals on critical thinking skills, AI governance, as well as improving processes and outputs. This requires us to build knowledge, skills and behaviors of how Busara staff interact with AI where it supports this vision, and maintain AI-free work where it does not.

By 2028, Busara is an organization of AI-literate critical thinkers, with 85% of existing staff and 100% of new staff having completed training for AI-literate critical thinking, after which 75% show increased AI literacy in testing that is delivered via AI with human oversight. A limited number of processes run with greater skill and efficiency:

a. *a curated knowledge management database, with AI-automated metadata tagging, reduces time spent searching for information by 40%;*
b. *automated horizon scanning delivers business development opportunities;*
c. *automated tagging of Busara mentions builds impact cases;*
d. *an AI-augmented dynamic tool on budget, bandwidth and likelihood of winning raises the quality of decisions on which requests for proposals to pursue. Technical quality, measured through qualitative internal and external feedback, has improved through deeper engagement with literature, supported by AI-driven literature scans validated against agreed quality standards. 25% of research projects contribute to a research agenda on the role of AI in social science research with and on the Global Majority. 100% of substantive outputs carry a Statement of Intellectual Fellowship and Accountability (sIfA).[6] As a result of having identified flawed processes for resource allocation, 90% of projects conclude on budget and on time. An audit of internal processes has produced a second round of AI use cases for beyond 2028.*

Key values and themes

Our vision is supported by taking a values-first approach to a general philosophy of innovative AI-in-the-loop configurations. Busara sees AI as a possible support of, or obstacle to, inclusive socio-economic development and equitable knowledge production. AI is for us not a tool to be used for efficiency, but an invitation to make us investigate efficiency as an ideological stance that we put in a relationship with innovation, quality, transparency and learning. As this touches Busara's core mission of using behavioral science to alleviate poverty, these values have to expand to include considerations of the role of AI in the following way:[7]

6 (Saleh and Schomerus 2026b; Schomerus 2026)
7 This aligns with Kenya's national AI vision, which highlights the need to use AI for 'inclusive socio-economic development' (Ministry of Information 2025).



Act with curiosity

We want to learn. We want to never stop asking the difficult questions in our work. We want to meet each other with openness and interest.

Acting with curiosity in our AI strategy means to:

1. *Nourish and challenge our cognitive abilities for critical thinking: analysis, attention, reflection, remembering, planning, evaluation, language processing, problem-solving, open mindedness, decision-making, and inference.*
2. *Protecting those cognitive abilities with clear cognitive and substantive guardrails.*
3. *Develop AI skills by doing the work, making mistakes, and improving.*
4. *Learn through responsible experimentation with AI in data collection and analysis.*
5. *Host a growing body of research on how AI can or cannot support ethical data collection and analysis.*
6. *Host a growing body of research on how AI can improve research processes to produce the most relevant insights for shaping policies and programs to support populations in contextually-appropriate and relevant ways.*



Act with respect

We want our work to be thoughtful and of high quality. We want to fulfill the need that we originally identified in pursuing a piece of work. We are careful about the complex issues we tackle, and mindful of the huge impact they have on people's real lives.

Acting with respect in our AI strategy means to:

1. *Commit to reflective and transparent AI use in all of our outputs, including by being at the forefront of global transparency standards.*
2. *Establish and hold ourselves to AI use quality standards that always serve our interest to improve the rigor and contextualized depth of our work.*
3. *Use the arrival of AI to build a more systematic Busara. This means to identify where processes and data care are flawed and how these can be improved (where helpful with AI, where based on behavior or lack of capability through culture change and training).*
4. *Use the arrival of AI to build 'one RAP way' that champions delivery, documentation, and reuse.*



Act Together

We are collaborators, seeking to actively share knowledge and perspectives. We seek to communicate and contribute to achieving our shared purpose.

Acting together in our AI strategy means:

1. *Work on trusting internal relationships and clear articulation and mitigation measures when trust breaks down.*
2. *Active, collaborative and responsible knowledge sharing for more visibility and use of our work.*
3. *Internal AI-supported knowledge management, including avoiding informal knowledge storing practices.*
4. *Delivering diverse outputs that bring a clearer Global Majority voice on responsible AI in behavioral science.*
5. *Championing diversity through different languages, different writing and speaking styles, different methods, different types of analysis, and different epistemologies. We produce diverse and contextually-rich work that represents not just Busara's multicultural team, but the many different places where we work, and actively works against a flattening of voice as driven by AI.*



Act with purpose

We want to keep a clear line of sight of why we do what we do. We want to take responsible ownership of our work, and be held accountable to the standards we set for ourselves.

Acting with purpose in our AI strategy means to:

1. *Take responsible ownership of our outputs.*
2. *Being transparent on AI use.*
3. *Highlighting the environmental costs of any AI usage.*
4. *Articulating trade-offs between short-term AI use and long-term impact on skills as well as equitable and sustainable socio-economic development.*
5. *Utilize AI to question our approaches and decisions that require alignment with our strategic goals and values.*
6. *Use AI to enforce substantive, ethical and cognitive guardrails.*

Operating principles

A number of operating principles emerge from these values:

- Knowledge management: Our work is more rigorous and innovative if we draw on more varied and better information.
- Administration: We can be innovative if our processes are of the highest quality.
- Data collection and analysis: Our work is more rigorous and innovative if our data is of the highest quality, collected ethically, and stored properly and our analysis is helped by challenging thought partners.
- Voice, impact and engagements: Writing is thinking, analysis and championing different voices, which is the core of what humans at Busara do.
- Networking and business development: Strong relationships are the core of what we do. AI supports us in building them by managing networking and responses.
- Decision-making support: Busara is driven by its values, mission and vision. We use AI to remind us of them, particularly when we are making decisions about using AI. We use it only to offer information to support better decisions.

Goals for the AI strategy

Goal 1: Build AI-literate, critical thinkers

Values: *Act with curiosity*

Goal-relevant use cases: Use case 1: Critical thinking coaching; Use case 2: AI-related project work inventory

- Specific: All staff develop AI literacy and critical thinking skills through structured training and hands-on experimentation, under guardrails that protect analysis, writing, and reasoning as human work.
- Busara develops its own body of research on AI's role in social science with the Global Majority.
- Measurable: 85% of all staff complete the AI-Literate Critical Thinker pathway (calibrated to role/seniority), each completing two supervised AI-analysis experiments with a documented reflection; 100% of onboarded staff complete the pathway; 75% show an AI-assessed (with human oversight) increase in AI literacy.
- 25% of research projects build the global knowledge base on AI in the empirical social sciences.
- Busara publishes a quarterly synthesis on AI and social science research with the Global Majority, building up to a major conference or publication milestone in 2028.
- Attractive: Builds staff confidence and judgment in using AI responsibly, protecting their cognitive skills to remain relevant and competitive. Elevates Busara's thought leadership at the intersection of AI and social science research.
- Realistic: Achievable through a phased training rollout combined with existing project work as the vehicle for analysis experiments. The research agenda can start now with work already scoped.
- Timebound: Q4 2026: Baseline data collection. Q1 2027: Design training program and research agenda. Q2 2027: Begin quarterly synthesis publications on AI in research. Q4 2027: Roll out training, reaching 85% completion by Q4 2028. Q3 2028: Major public conference or publication.
- Baseline needs (Q4 2026): Headcount of research/program staff and current self-reported AI literacy (skills audit or survey); a working rubric for how critical thinking is preserved to score reflections against over time; inventory of any existing experimentation with AI already happening across teams; inventory of existing AI-related projects.

Goal 2: Establish responsible AI governance

Values: Act with respect; Act with purpose

Goal-relevant use cases: Use case 3: Transparency and AI governance

- Specific: AI use in Busara is regulated through a shared platform and an evolving AI use policy. All outputs are fully transparent on AI use. All outputs pass an AI-supported quality check. All projects or proposals work with a published AI use policy and an AI impact statement covering disclosure, quality review, tradeoffs, skill building aspects, and environmental impact.
- All partnerships are checked against the partner's AI ethics-and-values for compatibility.
- Measurable: Busara works with an identified tool across the organization. Busara has a published and living AI use policy. 100% of substantive outputs carry a sIfA and pass quality review before release. All project proposals/reports include an AI impact disclosure. All partner/project decisions run through an AI ethics-and-values checklist. Busara publishes an annual summary of aggregate findings on AI in Busara in Tafakari.
- Attractive: Clear rules on how to work with AI and how to choose partners to counter uncertainty. Positions Busara at the forefront of global AI transparency standards and builds trust with partners, funders, and the populations the work serves.
- Realistic: Achievable by building on existing reporting templates, adding disclosure/review steps (including the CRediT taxonomy)/checklists rather than creating new processes from scratch. sIfA tool already exists and is starting to be used.
- Timebound: Q4 2026: tool/platform decision. Q4 2026: Draft standard and benchmark against global transparency norms, establish ethics- and values checklist, build do and don't list (AI use policy) for AI use. Q1 2027: Publish transparency standard and ethics-and-values checklist;

reach 100% compliance on new outputs; publish first annual public summary of AI in Busara. Q3 2027: decide whether pursuing an ISO 42001/IEC vs. NIST AI RMF approach is relevant for Busara.

- Baseline needs (Q4 2026): Audit of the last 6–12 months of external outputs: how many used AI, and was this disclosed? Draw a sample of past projects to retroactively estimate a year zero environmental/ trade-off baseline. Verify if any partner ethics-and-values checking processes exist. Benchmark against 2–3 global transparency standards or governance frameworks to calibrate the bar.

Goal 3: Identify and operationalize AI opportunities for 'one RAP way' for systematic delivery

Values: Act with respect

Goal-relevant use cases: Use case 4: Meeting notes and action-item tracking; Use case 5: Project management; Use case 6: Building a queryable institutional memory that enforces data protection and prioritizes AI on the input side; Use case 8: AI-assisted public adoption alerts; Use case 10: Proposal/ BD support

- Specific: All research, advisory or practice projects work from a single template (using Busara's AUDAS model) standardizing kick-offs, version control, documentation, and closures.
- A common RAP language protects a consistent Busara identity across staff and client transitions.
- Measurable: AI-supported automation is introduced for: budget forecasting (including dynamic budgeting/contracting based on scenarios), project and time management/tracking, proposal sense-checking, maintenance of proposal success-rate statistics, reaching full adoption by Q3 2028.
- 100% of active projects work 'one RAP way', reducing average project setup/handover time.

- Adoption alerts show an upward trend. AI assisted public alerts (e.g., partner publications, strategy documents and program announcements) help surface possible evidence of adoption that is always checked with human validation.
- Attractive: Reduces duplication of effort, protects institutional knowledge when staff transition, and raises the consistency and reusability of the organization's analytical work; eases workload on administrative and financial processes.
- Realistic: Achievable by consolidating existing workflows into one template, piloting automation on the administrative processes with the clearest efficiency case (e.g., budget forecasting, proposal win rate tracking) before wider rollout, and providing training/support during rollout and using AI to supervise compliance.
- Timebound: Q4 2026: Build baseline of adoption alerts. Q1 2027: Identify and design the standard RAP template. Q2 2027: Pilot administrative/financial automations (budget forecasting, proposal sense-checking). Q4 2027: Roll out to all active projects; full adoption and time reduction measured by Q3 2028. Q1 2028: check adoption alerts against baseline.
- Baseline needs (Q4 2026): Current number of distinct analysis workflows/templates in use across teams; Average time currently spent on project setup, handover, or redoing lost/undocumented work; Existing code/documentation reuse rate (% of projects starting from a prior template vs. from scratch).

Goal 4: Utilize AI to improve delivery of diverse, high-quality, and operationally-relevant publications

Values: *Act together; Act with curiosity*

Goal-relevant use cases: Use case 6: Building a queryable institutional memory that enforces data protection and prioritizes AI on the input side; Use case 7: Tagging/classification support for the outputs map; Use case 8: AI-assisted public adoption alerts; Use case 9: First-pass literature screening and lit-review support

- Specific: Institutional knowledge is centralized and searchable, supporting proposals, evidence synthesis, and better outputs. Projects produce outputs that accurately represent researched communities' voices. External publications are diverse, high-quality, and operationally-relevant, fostering partnership and business-development opportunities. Outputs have operational relevance. Outputs pass a designated review step to check for bias. Busara does not publish work that is largely AI written.
- Measurable: Metadata and a rigor score are co-created for better knowledge management. 100% of highly-relevant institutional knowledge is migrated into a shared, AI-searchable database (with static and Folder sync elements), with an archived database for less relevant documents. At least four outputs per year amplify Global Majority perspectives on responsible AI. Every project produces a project summary (Aha! Moment) and one audience-targeted output. 70% of surveyed readers rate publications as clear and usable for decision-making.
- Attractive: Counters frustrations caused by messy knowledge management leading to duplication of work or lower quality when work is produced under time pressure without access to institutional knowledge and instead strengthens learning and presentation of past work. Facilitates research being used in operational decisions in contextually appropriate ways. Champions diverse voices and reduces the flattening effect of AI on writing and analysis styles. Frees up staff time to focus on substantive quality (better knowledge outputs, outreach, and creative engagement design).
- Realistic: Achievable by building on existing publication workflows and knowledge-sharing platforms, with incremental migration and format changes. Audience profiling and multi-format generation can build incrementally on the knowledge base once established.
- Timebound: Q1 2027: audit knowledge storage, design metadata and rigor score, build database. Q2 2027: begin migration and training. Q1 2028: reach 90% migration; survey relevant audiences for their

needs. Q3 2028: design and trial audience-profiled products; reach 100% migration. Q4 2028: 80% of projects have practitioner-facing publications.

- Baseline needs (Q4 2026): Current publication formats: how many already have a practitioner-facing version vs. General-readership-only; Survey of past readers/users on clarity, usability, and whether any past publication was actually applied in a decision; Conducted for this strategy or pre-existing: Staff time currently lost searching for or recreating existing work (proxy for knowledge fragmentation cost); Full audit of where knowledge currently lives (% on shared drives vs. informal/personal storage); Count of existing external outputs in the past year and who authored/voiced them.

Goal 5: Diagnose where AI can help fix flawed processes and data practices

Values: *Act with respect*

Goal-relevant use cases: Not applicable

- Specific: Flaws in core RAP, Research Excellence, BizOps and FIVE processes and data practices are systematically catalogued to capture why projects miss deadlines, budgets, quality or impact to identify further use cases where AI might close a genuine gap or where culture or behavior change is a priority.
- Measurable: Complete a full audit of data processes, classifying flaws by root cause. Resolve 70% of high-priority flaws. Achieve a 50% reduction in data-quality incidents. 90% of projects delivered within time and resource budgets. A catalogue of AI use cases beyond 2028 is developed.
- Attractive: Builds a more systematic, rigorous Busara, reducing wasted effort and improving the reliability of the organization's data and research outputs.

- Realistic: Achievable by first mapping the pipeline and classifying flaws before committing to specific fixes, ensuring resources go to the right type of intervention. Realistic because it works in close connection with goals 3 and 4.
- Timebound: Q2 2027: Map current processes (including building a baseline on data-quality incidents) to flag repetitive tasks, pain points and causes of project time and budget overrun. Q4 2027: Complete audit and classify flaws by root cause. Implement fixes; reach 70% resolution and 50% incident reduction by Q1 2028. Q4 2028: 90% of projects delivered on time and budget. Identify future use cases.
- Baseline needs (Q2 2027): Full process audit mapping current processes, flagging known pain points (through platform experimentation and staff survey); Classification of existing flaws by root cause (technical/tooling vs. behavioral/cultural), to determine whether AI or training is the right fix; Historical rate of data quality incidents (errors, rework, inconsistent coding) over the past 12 months.

Goal 6: Utilize AI to improve networking and business development

Values: *Act with purpose; Act together*

Goal-relevant use cases: Use case 6: Building a queryable institutional memory that enforces data protection and prioritizes AI on the input side; Use case 10: Proposal/BD support

- Specific: AI agents support Busara in scanning and assessing opportunities, identifying networking contacts, and responding swiftly to proposals with set pieces drawn from Busara's knowledge base — easing a major workload bottleneck while keeping human judgment central.
- Measurable: 80% of pursued opportunities/proposals that respond to RFPs (meaning they are not part of the partnership strategy) originate from automated scanning. Set-piece pre-population is available for

all standard proposal types. 90% of proposals query the internal knowledge database. Every AI-suggested opportunity passes the Goal 2 ethics-and-values checklist. 75% of conference attendees use the Busara Conference networking tool.

- Attractive: Eases a major workload bottleneck through more targeted, faster proposal submissions, and strengthens Busara's visibility and network reach.
- Realistic: Depends on Goal 4's knowledge management system, since opportunity matching and set-piece generation rely on a queryable knowledge base of Busara's past work and portfolio. Depends on Goal 3's identification of processes.
- Timebound: Q3 2026: refine Busara Conference networking tool. Q4 2026: design horizon scanning tool. Q3 2027: Develop opportunity-scanning and set-piece system (building on Goal 4's knowledge base). Q1 2028: System fully operational; measure against 80% target.
- Baseline needs (Q1 2027): current proposal success rate/ based on relationships/ level of input/ referrals not holistically tracked.

Transformation, partner experience and market position by 2028

Busara staff will meet AI as highly-skilled critical thinkers, using it to improve quality and innovation rather than efficiency alone. Processes such as horizon scanning and adoption alerts run automatically with humans in the loop. Knowledge management is designed as AI-in-the-loop, alongside stronger human skills in storing, querying and using knowledge. AI serves human-initiated quality checks in both research outputs and managerial processes.

We become a stronger partner through greater trust in our work, helped by our radical honesty on transparency, global thought leadership on AI governance, and catalogued learning on the merits and limits of conducting social science research with AI support.

We distinguish ourselves from competitors through our holistic and thoughtful approach to AI (we are not suddenly an AI-first organization); our ethical commitments in research with Global Majority populations, some of whom may be tech-illiterate; through continuous assessment of AI's impact on those populations and on the research conducted with them; and by publishing our learning as a public good with radical honesty about AI use. We continue to train and safeguard staff cognitive abilities for their AI-in-the-loop role. This is markedly different from the efficiency or AI-first framing most competitors are adopting.

Measurement

Below, Table 2 shows the strategic coherence between the org and the AI strategies. While this table focuses on numerical KPIs, it is important to also offer narrative assessments of what works.

Table 2. How do org KPIs map onto AI strategy KPIs?

Org KPI	AI-Strategy Goals	AI Strategy KPIs
KPI 1: Total revenue. Are partners choosing Busara at the level needed to sustain the organization?	**Goal 3**	Proposal sense-checking; maintenance of proposal success-rate statistics (full adoption by Q3 2028)
	Goal 6	80% of pursued opportunities/ proposals originate from automated scanning. Set-piece pre-population is available for all standard proposal types

Org KPI	AI-Strategy Goals	AI Strategy KPIs
KPI 1: Total revenue. Are partners choosing Busara at the level needed to sustain the organization?	Goal 3	Every AI-suggested opportunity passes the Goal 2 ethics-and-values checklist
	Goal 6	90% of proposals query the internal knowledge database
	Goal 6	75% of conference attendees use the Busara Conference networking tool
KPI 2: Gross profit margin. Are we delivering that work sustainably and efficiently?	Goal 3	Budget forecasting (including dynamic budgeting/contracting based on scenarios); project and time management/tracking (full adoption by Q3 2028)
	Goal 3	100% of active projects work 'one RAP way', reducing average project setup/handover time
	Goal 5	90% of projects delivered within time and resource budgets
	Goal 5	Resolve 70% of high-priority flaws
	Goal 5	Achieve a 50% reduction in data-quality incidents
KPI 3: Practice partnership share. Is the portfolio moving toward the strategy's headline bet — ≥60% revenue from renewable, practice-aligned partnerships by 2028?	Goal 2	All partner/project decisions run through an AI ethics-and-values checklist

Org KPI	AI-Strategy Goals	AI Strategy KPIs
KPI 4: Scaled adoption rate. Is our work taking root in partner systems beyond pilots?	Goal 3	Adoption alerts show an upward trend. AI assisted public alerts (e.g., partner publications, strategy documents and program announcements) to support Google citations help surface possible evidence of adoption (checked with human validation)
KPI 5: Evidence and learning asset rate. Are projects turning into reusable knowledge, tools, and assets?	Goal 1	25% of research projects build the global knowledge base on AI in the empirical social sciences
	Goal 1	Busara publishes a quarterly synthesis on AI and social science research with the Global Majority, building up to a major conference or publication milestone in 2028
	Goal 4	Metadata and a rigor score are co-created for better knowledge management
	Goal 4	100% of highly-relevant institutional knowledge is migrated into a shared, AI-searchable database (with static and Folder sync elements), with an archived database for less relevant documents
	Goal 4	At least four outputs per year amplify Global Majority perspectives on responsible AI
	Goal 4	Every project produces a project summary (Aha! Moment) and one audience-targeted output
	Goal 3	70% of surveyed readers rate publications as clear and usable for decision-making

Org KPI	AI-Strategy Goals	AI Strategy KPIs
Cross-cutting	Goal 1	85% of all staff complete the AI-Literate Critical Thinker pathway (calibrated to role/seniority), each completing two supervised AI-analysis experiments with a documented reflection; 100% of onboarded staff complete the pathway; 75% show an AI-assessed (with human oversight) increase in AI literacy
	Goal 2	Busara works with an identified tool across the organization
	Goal 2	Busara has a published and living AI use policy
	Goal 2	100% of substantive outputs carry a sIfA and pass quality review before release
	Goal 2	All project proposals/reports include an AI impact disclosure
	Goal 2	Busara publishes an annual summary of aggregate findings on AI in Busara in Tafakari
	Goal 5	Complete a full audit of data processes, classifying flaws by root cause
	Goal 5	An audit of internal processes, completed in H2 2028, identifies the second round of AI use cases for beyond 2028

A closing review in Q4 2028 assesses the AI strategy against all five org-KPIs, combining the numerical measures above with a narrative assessment. Its findings feed the audit of internal processes that identifies the second round of AI use cases for beyond 2028.

Value proposition of AI

Calculating an exact value proposition is challenging in the context of Busara: generally, opportunity cost is higher than cash outlay, meaning poor processes mean we miss out on opportunities for better work. Overall, the value proposition focuses on value-based improvements through better research infrastructure.

- For Busara, low-threshold access tiered to different levels of AI skill supports equitable upskilling, streamlines and improves processes, opens access to knowledge, and carries set-piece work so innovation can happen where it matters. Because growth can be a consequence of impact, this approach could increase the evidence and learning asset rate and total revenue, contributing to org-KPIs.
- For our partners, the AI strategy creates a Busara that meets the moment: anticipating change with innovation, and resolving long-standing issues of governance, quality, skills, values and processes while delivering more reliably and sustainably on its core work of human knowledge creation, contributing to org-KPIs.
- The broader societal benefit lies in sharing our learning and governance of AI for ethical research with the Global Majority as a common good, to improve global practices and outcomes for all.

Industry-specific applications

Knowledge creation within international development is a distinct sector for AI purposes. Our product is ethical, trustworthy, and operationally-relevant knowledge; this often involves conducting research with marginalized communities, some of whom might be illiterate or tech-illiterate. In addition, much of our research can be shaped by funder interests, which are not always aligned with our own interests or what affected communities consider relevant. This means our approach to AI adoption is truly not a technical one: we treat this as a moment in which social science skills matter. We analyze, we look at impact holistically, and we understand that our decisions are influenced by our surroundings. This means to prioritize AI governance, equitable skill building, and transparency while acknowledging that internally, we are still building our own institution's capacity, capability, processes and compliance.

Research integrity and disclosure are sector-defining concerns

- Ethical knowledge production is our sector's product, meaning that any use of AI is not a question of efficiency only, but of integrity and values. Reasserting this stance counters the urge to use AI for speeding up work, which also is the source of Busara's 'dirty secret' logic of AI use, which this strategy seeks to counter comprehensively. Reversing this and setting a trend of transparency and values first is a matter for our organization and for the research sector more broadly.
- While there is currently no single, agreed standard for how AI use should be disclosed in research outputs, this is a vibrant topic of debate. International actors and networks, such as the World Conference on Research Integrity Foundation, the International Science Council, COPE, STM, Global Young Academy, as well as actors in which Busara is involved (such as Behavioural Research UK) are actively putting forward. recommendations for standards, including seeking consensus

for a Global Reporting Standard for AI Disclosure in Research. Busara, with its committed stance on transparency and common good contribution in the shape of the sIfA (with the accompanying background paper on where the sIfA's philosophy sits within global scholarship on the matter) can and must participate in this active process.[8]

- Busara's disclosure standard as reasserting values and bold steps. Busara's sIfA extends the widely-used CRediT taxonomy of 14 standard research contributor roles to also capture AI's contribution per role, on a 0–2 extent scale, with reasons and named tools. sIfA gives Busara a sector-appropriate, designed-and-shipped answer to how we disclose AI use, rather than needing to adopt someone else's generic policy (which speaks to one of the five bold steps on writing the rules on AI governance for ourselves). The sIfA is designed to be kept as a project log alongside all steps and contributors, offering a simple and time-efficient mechanism.
- Maintaining value-driven parts of our work as entirely AI free (such as writing as a way to protect cognitive abilities or not flatten Busara's voice) and disclose this firms up our position as a value-driven learning organization that is prepared to take a committed stance and thought leadership.

In research and behavioral science, language matters

- Because Busara trades in credibility, interpretation and championing different voices, it matters tremendously that we talk about AI in thoughtful and context-specific ways. One suggestion is to counter anthropomorphizing language in line with the recommendations of Bender and Inie in how we use language, describing technology by function ('automatic transcription' or 'text generation') rather than cognitive terms (writing, interpreting, analyzing), clarify that agency

8 (Saleh and Schomerus 2026b)

lies actively with the human (no passive language that explains that 'an LLM was used') and name tools, rather than generic AI. Adopting a reflective stance on the use of AI language as part of our style guide could become a differentiator for Busara.[9]

Busara can differentiate itself as a Global Majority representative, but differently than in the past

- Busara can and must take the position of a spokesperson and thought leader on Global Majority representation in AI systems and technology development. While this is a subset of larger questions about the role of AI in the social sciences, it allows for a more specific claim and unique proposition that ought to be central to what we do. This builds on previous work that Busara does to address structural problems in North/South research collaborations, with questions of credit, visibility, voice and knowledge hierarchies continuing to be a challenge. Using sIfA strengthens the additional AI angle in our existing commitment to knowledge equity.
- Busara has been in the room when sector-defining outputs were discussed, such as the CGD/Agency Fund GenAI Evaluation Playbook. Our stance is that such frameworks often miss points of participant data agency, take a narrow view of qualitative research, and do not account for learning through reflection and not just indicators alone. It will be worthwhile for us to continue articulating those stances and publicize them further as a differentiator to our competitors.[10]

Contractual AI restrictions can vary vastly

- Only one of Busara's contracts imposes strict no-AI policies, which can create friction with our current situation of shadow AI. The strategy is

9 (Bender and Inie 2026 (7 January))
10 (Center for Global Development, The Agency Fund, and IDInsight 2026)

designed to have a strong value base from which all client requirements can be met or demands that overstretch our values can be pushed back (and might become part of the partner suitability check list). Practically speaking this requires an extra level of contractual negotiations and checking to make sure we are not put in a position where we cannot fulfill a too-wide AI mandate.

Different regulations apply to our work

- We are guided by a number of regulatory frameworks in addition to our own standards and values. These require careful monitoring: Kenya's pending AI Bill identified some Busara-relevant sectors as high risk (healthcare, education, agriculture, finance, security, employment, and public administration, under which Busara's work on governance and conflict would fall). This means vigilance to monitor if any of Busara's work will fall into a high-risk regulatory tier and how we adjust for that.
- At the same time, Kenya's national AI Strategy (2025–2030) has named Busara-relevant sectors as priorities (healthcare, agriculture, and public service delivery), which corresponds to two of Busara's identified practice areas. This external validation ought to serve as a fundraising energizer if utilized explicitly and as an impact factor when seeking to influence government policy under this framing.

Busara's market context

Client and funder landscape

Busara's top funders are philanthropic foundations, governments, multilateral organizations (UN), academic and research institutions, and the private sector. Contractual treatment of AI use is extremely limited, with the issue often discussed informally rather than contractually. No contract yet includes an AI disclosure agreement, and few requests for proposals stipulate exact rules on AI use.

Some requests for proposals mention AI explicitly. There are three versions of this mention: the request to explicitly use AI (and disclose the organization's AI capabilities), to explicitly disclose AI usage, or to avoid AI altogether. Currently, Busara does not have standard procedures on how to respond credibly and transparently to each request.

Data protection policies are established and stretch across all jurisdictions in which Busara works. Contractual clauses on how respondent data is handled are clear, guided by compliance laws in the relevant jurisdictions, and pursue compliance both with the letter and spirit of the law. In practice, this means that data sharing with clients would mean that the client has institutional review board (IRB) approval, a signed data-sharing agreement is in place (including safeguards for cross-border transfer), and identifiers have been removed beforehand. Participants also need to be told in the consent form how their data may be shared and used, and to consent to that specifically. However, even when legally data sharing is possible (which currently it is not, due to consent), considerations on the ethics of it will always be made on a case-by-case basis.

What are Busara's peer organizations doing with AI?

Broadly speaking, Busara's peer organizations are taking a very different approach to engaging with AI compared to Busara's as outlined in this strategy, as Table 3 next illustrates.

Table 3. Broad comparison of how Busara's peers approach their AI strategy?

Busara's peers and competitors	Busara
Tool and service-driven: how are we using AI?	Value-driven and fundamental: how does AI connect to our values; what do we need to take into account in how we fundamentally do our work?
Many organizations relaunched as AI-native or AI-first	Busara's approach is rooted in organizational context and realities
Others asserting their readiness to working with AI	Busara is being honest about readiness. We need to first fix our processes fully rather than automating poor processes. We also need to build AI capacity first
Varying degrees of transparency on AI use	Radical honesty on AI transparency

A number of relevant organizations have published AI transformation benchmarks and lessons learned that must inform Busara's approach.

IDInsight's internal AI transformation experience (three years in, 98% weekly generative-AI usage among staff) names four common pitfalls from which Busara can learn:

- Option paralysis: The recommendation is to not spend months evaluating tools, but start with one enterprise-grade platform and configure it. This recommendation has informed Busara's approach.
- Lack of governance/safety clarity: Despite taking a longer, more reflective approach, an immediate and regularly updated reference sheet of do/don't do is a necessity to keep the learning alive. This recommendation is directly reflected in Busara's ethos to offer guardrails swiftly (building on the SCP's PROTECT framework) and then

pursue an open, values-driven approach rather than a one-off policing approach to keeping governance alive.

- No time to learn: IDInsight stresses the need for building staff capacity through short and accessible learning formats (how-to videos, peer-shared use cases in team meetings) and offer dedicated time to internal AI values ambassadors, rather than hoping for a general dedication to become autodidacts. This recommendation is mirrored in Busara's change management plan.
- Not measuring progress: It is easy for AI use to plateau on low-leverage tasks that are mainly driven by efficiency thinking (such as meeting summaries). By establishing measurement that captures higher AI ambitions, use cases can become more sophisticated and offer higher return, particularly when it comes to improving quality and innovation. This recommendation is directly reflected in Busara's Goal 1 to systematically build AI-literate critical thinkers.[11]
- Organizations seeking to establish what ethical AI use looks like in the research/development/humanitarian sector are offering relevant governance benchmarks. BenevAI Certification, for example, focuses on governance, risk management, data lifecycle, accountability, and transparency. The Centre for Governance Innovation Human Analysis Standard, along with BR-UK's declaration also offer similar check points for Busara when finalizing our own ethics and security strategies.[12] Busara is in conversation with these actors, primarily to include considerations of knowledge creation in the BenevAI certification and to include sIfA in the tools that can be used for AI transparency.[13]
- Research-disclosure and academic-integrity benchmarks: With sIfA, Busara already follows best practices of the CRediT taxonomy in declaring contributions to research outputs. sIfA also adopts existing best practices put forward by publishers on maintaining human accountability for all research output.

11 (Sampson and Singh Rawat 2025)
12 (Centre for International Governance Innovation 2026) (Behavioural Research UK 2026)
13 www.benevai.org

Stakeholders and change management

The stakeholder and change management model recognizes the diversity of internal (differentiated by role, departments and AI archetypes) and external stakeholders (differentiated by those funding research, using research or participating in research) and that many of these have unclear interests due to the fluid nature of AI use and development. An accompanying change management concept (designated staff, broad-based project involvement, numerous communication channels to support diverse learning and communication styles) supports engagement with the AI strategy, the decisions that need to be made, and with the outcome.

The AI strategy has two stakeholder assumptions: that all departments in Busara will have to be part of the organization's AI governance and skill set, and that our approach to knowledge equity requires that all members of staff have access to the AI Literate Critical Thinker pathway.

Busara's culture and attitudes towards AI

Because no strategy will take hold if it is alien to the culture of an organization and within teams, it is necessary to unpack concerns, hopes and fears for AI in Busara. This section draws from the consultations, surveys and exec team meetings (for detail on the methodology, see 16.5 Appendix 5: Methods used in developing this AI strategy.

Six phases of org-wide consultations (March through August 2026) tracked a shift from silence to openness. Early sessions surfaced AI use as a 'dirty secret in Busara; used but undisclosed. By April, staff named concrete use cases (coding, risk scoping, reviewing health results) and voiced hopes and fears for their own thinking. This disclosure gap prompted Busara Futures to design and ship the open-source sIfA tool, which uses existing community standards

to declare human/AI interaction in knowledge production (www.sifatool.org). Through June, emotional responses ranged from AI enthusiasm to shame over labor exploitation in model training. By July, consultations shifted to co-designing solutions. Separate executive team sessions raised a distinct concern: uneven adoption, not AI itself, is the primary organizational risk.

The different modes of consultations surfaced different priorities. While deskilling was a major topic of discussion in the consultations, it ranked only fourth of concerns in the survey (61%, against 82% for accuracy/reliability, 80% for over-reliance, and 74% for data privacy). Data privacy ranks high as a concern in the survey, but barely comes up in discussion. While administrative and documentation tasks were cited most in the survey as the tasks staff would most hope for AI to help with (51%, the highest of any listed task), they were not voted an area of interest in the consultations, possibly reflecting who attended the consultations.[14]

Conversations on AI as part of this strategy development started in early March, where the overall tone was described in a Claude Sonnet 5 sentiment analysis of transcripts as 'abstract and guarded', with philosophical thoughts about what AI is in the organization dominant, for example the "AK-47s people don't know how to use." During this set of consultations, questions of value systems dominated and clashed directly with operational pressures on revenue, delivering proposals in minimal timeframes, and a general guardedness when it comes to individual learning in the organization, most evident by the widespread culture that people do not share first drafts (and that these first drafts often become final drafts still stored on people's private drives, adding to the knowledge management problem).

The second phase of consultations in late March 2026 started by explicitly unearthing the dirty secret, which shifted the tone towards more openness. People named specific workflows for which they use AI (coding, risk scoping, reading of health results). We talked about personal hopes and fears connected to AI: what will it do to our minds? Our cognitive abilities?

14 (Saleh and Schomerus 2026a)

Based on the insights that Busara's 'dirty secret' and the need to build a culture of openness and transparency (in itself a major behavioral and culture change) was not unique to the organization, but in fact a global problem, the two leaders of Busara Futures developed the sIfA tool. During a specially-convened consultation with the MRI Team (phase three of the consultations, April 2026), the debate focused on the specifics of the sIfA tool: is its default setting to the CRediT framework helpful? What scale should AI contribution be measured on? Is the depiction of a figure alongside a table helpful or distracting? Does AI disclosure equate support for using AI? Is scoring AI use a useful way forward? Some of the feedback shaped the design of the sIfA tool (such as the measurement scale), while areas of disagreement were addressed in the accompanying Groundwork.[15]

Phase four throughout June marked the org-wide consultations on what it means to have an AI strategy and AI vision. Attendance was acceptable, but on the low side when looked at it from an organizational perspective. Checking in on what emotions people have about AI, the register ranged from enthusiastic; dismayed, specifically about junior staff being 'cheated out of opportunities to struggle through problems that really build character'; astonished/awestruck; enthralled; shameful due to the labor exploitation in model training that underpins AI; contempt; agony; helplessness; eagerness; opportunities, including around self-discovery.

Some tension emerged about the approach the AI strategy consultations were taking—for example, when the approach explicitly focused on pain points or how possible time saved could best be spent, rather than on eliciting more generic opinions about AI. Additionally, questions emerged on how Busara could tout innovation and quality in its approach to AI when revenue pressures consistently foreground efficiency.

Phase five in July focused on co-designing what AI solutions could look like. Suggestions included translation-consistency lookup, active anti-duplication

15 (Saleh and Schomerus 2026b)

of survey instruments, a rigor-scoring mechanism, formatting/syntax checks, precision extraction of limitations sections (the specific use cases are further elaborated in the section Solutions and resources).

In parallel to the org-wide consultations, the executive team had its own separate learning and aligning sessions (three in total) including asking to participate in a written strategic exercise on pain points. Some concerns mentioned here did not appear in org-wide conversations: that not AI per se poses a risk to Busara, but uneven adoption does; that AI could support evening out inconsistent quality standards across portfolios, or help scale a unified 'one RAP Way' delivery methodology.

What—in theory—do Busarians value as Busara culture?

Five points repeatedly come up as values that should (and sometimes do) guide Busara's decisions and daily practices.

- Thinking: this value emphasizes the importance of challenging thinking, rigorous knowledge and the ability to deal with contrasting views.
- Different voices: this value emphasizes the need to produce diverse and contextually-rich work that represents not just Busara's multicultural team, but the many different places where the work is done. This value extends to championing different writing and speaking styles.
- Openness to challenge: Substantive disagreement, rooted in research, is welcome, but must be articulated rigorously, rather than expressed as a discomfort.
- Rigor and depth: Rigor and contextualized depth of the work are high aspirations for Busara and simultaneously two of the greatest pain points, with much work considered as lacking both.
- Learning: As a value, the organization champions giving people the opportunity to develop skills by doing the work, making mistakes, and improving.

Box 1: The spirit of what Busara seeks in an AI strategy

Across the various consultations, a number of themes emerged of what Busarians expect from an AI strategy:

No forced AI usage: the option to remain AI-free;

Not simply policing of AI use, but constructive approaches driven by ethos and values;

Thoughtful AI policies that build existing skills and add new ones;

Finding a way to counter high levels of distrust of colleagues and leadership.

What is the sentiment about AI in Busara? Insights from the consultations

Stakeholder archetypes

- Broadly, attitudes within Busara fall into three camps:
- Learn and compete: let's use AI responsibly and thoughtfully to improve our work and secure our organizational survival.
- Fix the basics first: let's not let AI distract us from needing to fix our processes, structures and compliance first.
- No role for AI: let's not use AI as it has no place in knowledge production.

Within these camps, attitudes can be grouped into different archetypes. These archetypes developed below combine the answers of different people under archetypes. One individual's answers can show up in several archetypes. No archetype represents a single person. These archetypes can inform more nu-

anced configurations for AI tools and a more responsive change management process.

The Process Practitioner: The process practitioner understands that processes are currently lacking and seeks concrete practical use cases that work in supporting better processes, such as when using Claude to support making an organizational go/no go decision on a proposal or a project path or on automating timesheet processing. The process practitioner sees AI as a way to institute institutional knowledge by using AI to make processes replicable across teams and client transitions. This perspective thinks about using guardrails as part of performance management: in addition to a catalogue of org-wide guardrails, the suggestion is to assess an individual's guard rails. The process practitioner has personal guard rails already in place, with clear distinctions between acceptable and unacceptable use of AI (for example: writing run-of-the-mill code is fine, but writing code that the process practitioner could not write themselves is not). The process practitioner suggests that AI could include a guardrail for cognitive cooldown—meaning that AI will pause or query before answering high-stakes questions. The process practitioner wonders how to best train a Busara-specific AI platform, considering that guard rail needs can be diverse across the organization and Busara's corpus of work can be of highly uneven quality: should inclusion in the knowledge corpus be determined with a date cut off, through a peer-rated rigor score or a ranking exercise on what the best outputs are attached to knowledge-base outputs, or should only what has been officially published be included? Other questions that the process practitioner raises are about prioritizing longer-term enhancement of processes with AI, rather than thinking of AI as a one-time fix of a long-standing problem (such as the enormous knowledge management backlog). The process practitioner is skeptical about how useful established processes are in the age of AI—for example on using the CRediT framework for AI disclosure, arguing that the CRediT framework in itself has failed to create clarity on author contribution.

The Thinker: The thinker is longing for an organizational set up that incentivizes and nourishes thinking over emphasizing revenue creation only. The

thinker also sees the possibility that AI might ease the cognitive load created by project pressures to free up time for thinking.

The Futurist: The futurist imagines what today's behavior might mean for the future, including rapid deskilling because young staff members skip the struggles and exercises that help them build both their technical as well as character skills. Outsourcing the mundane tasks means that the mundane tasks can no longer support skill-building of those just starting to build a skill. The futurist worries that the modus of 'explain like I'm five' will erode learning and thinking.

The Diversifier: The diversifier worries about the homogenizing effect of AI, based on its bias and output style. The diversifier emphasizes the need for Busara to maintain its distinct voices and sharp awareness that bias can show up in subtle ways. The diversifier highlights that AI-supported outputs can be too general or too removed from local contexts, particularly when it comes to local knowledge from African contexts.

The Coworker: The coworker accepts AI as a helpful colleague, but one that is ideally always one step behind. In practice, this means that the guardrail ensures that AI defers to the coworker when the coworker has obvious strengths or already mastered a task: 'I want AI to leverage on my strengths, as opposed to me leveraging on its strengths'.

The Quality Controller: The quality controller sees AI's strength in checking on work that was done, and doing so consistently across the organization through a shared subscription to a model. The quality controller also acknowledges that better tools can support better quality. More broadly, the quality controller sees an opportunity to use AI to counter Busara's incentives to prioritize speed over rigor without incentives to meet agreed quality standards.

The Skeptic: The skeptic suggests that org-wide guardrails cannot work because too many different cases and jobs would need to be covered; that pro-

cedural problems in Busara are solvable without AI through better systems and enforcement; that declaring AI use means exerting implicit pressure to use AI, leading to alienation of those who do not want to; and that leadership in itself is not transparent about AI use. The skeptic assumes that guardrails cannot work as people will simply use private AI channels to circumvent them. The skeptic does not believe that there is a way to introduce AI fairly because once introduced, it leaves no options to opt out because everyone's work efficiency will be assessed on the basis of AI-enhanced efficiency. This will make entry-level jobs even more difficult. The skeptic worries that the language used to describe AI humanizes it. The skeptic does not accept the framing of AI as a challenging thought partner, highlighting that that role needs to fall to qualified colleagues and that AI might shortcut that perceived need to talk to others.

The Learner: The learner uses AI as a coach, for example through a self-curated course on relevant behavioral concepts such as COM-B.[16] The learner also interacts with AI more on a coaching level, pushing back against outputs that are inappropriate, deskilling or limit agency. The learner also highlights that training will need to continue to be a Busara priority, rather than being outsourced to AI.

The Philosopher: The philosopher asks what the arrival of AI means for humanity and for the research work in general.

The Principalist: The principalist seeks to establish org-wide principles to even out that AI is used by some creatively and effectively, and not at all by others.

The Ethicist: The ethicist foregrounds the human, environmental and cognitive costs and sees AI usage as endorsing unethical practices, for example, exploitation of Kenyan data workers. The ethicist also highlights that no amount of AI disclosure will give credit to the people who were used to train the models (such as authors and data curators).

16 (West and Michie 2021)

The Pragmatist: The pragmatist supports finding processes to deal with the current unclear situation to develop approaches that can then evolve. This includes using the sIfA tool to counter reputational risks of AI-non disclosure, and making the tool as practical as possible. It also includes finding answers to the questions of whether disclosing AI use might lead clients to cut budgets or delivery timelines. The pragmatist asks if AI can be used to offload rants, thus easing management of emotions.

The Common Gooder: The common gooder identifies the potential for Busara to contribute to AI as a better common good, for example, by getting involved in safely contributing to better models.

Hopes and fears for AI at Busara

When discussing what people hoped that AI could bring to their work, claiming back time for higher-value work is most common. Higher value work here is either thinking, learning, more teamwork and communal brainstorming (and fighting personal isolation through it) and focusing on well-being. Another hope is that AI can act as a democratizing learning tool, as AI support can make up for different levels of education and expand vocabulary. AI is also seen as an opportunity for Busara to make a public contribution by publicly sharing our own learning and becoming a trusted voice on what responsible AI looks like in behavioral science and for the Global Majority. A focus for that line of thinking is that Busara is well-versed in what it means to understand and examine context and lived experience and that this might not only continue to set Busara apart, but also make context setting for AI much better, accelerating the quality of AI-supported work.

Another hope is connected to the opportunity that AI brings to changing Busara's transparency and learning culture.



39% of people responded in the survey that they want'opportunities to experiment safely and engage in ongoing discussions', see the opportunity to learn through 'dedicated meetings / office hours' (53%) or 'webinars with internal and external speakers' (68%)[17]

Fears focus on deskilling, which includes not developing skills for junior staff or neglecting and thus losing skills already built. Organizationally, the fear is that these early days of AI could put the organization on a trajectory of skills loss that will make Busara irrelevant. Additionally, different levels of AI adoption and deskilling fuel the fear of unfairness, meaning that it is assumed that those who adopt AI might fare better in their career trajectory based on the presumed speed at which they work. Other fears are focused on the relationship between humans and technology, asking what it means to be human or how to talk about machines without linguistically injecting humanity in them. In Busara's offering, reputational concerns about using AI at all or about not being transparent about AI usage dominate, alongside the fear of smaller budgets and inconsistent quality.

Box 2: Busara was built to counter bias. Yet, AI bias is not a major topic

Despite Busara's emphasis on bias as a concern in behavioral research, the topic of AI bias only came up a few times—deskilling, data protection or reputational risks were the concerns more regularly noted. This raises the need for Busara to continue working on bias, including on staff awareness, a broader research agenda on what biased AI in the Global South looks like, and on how to mitigate bias when using any AI tools.

17 (Saleh and Schomerus 2026a)

Constraints and tensions

For the AI strategy to work, it is necessary to speak openly about constraints. The strategy must reckon with limited policy compliance, diffuse decision-making authority and low AI literacy in Busara.

Diffuse decision-making structures blur the question who has the authority to set any rules. Currently, AI literacy is low: for example, while most AI users report doing so ethically, AI disclosure is almost non-existent (a major ethical challenge).

Many structural issues exist, often due to non-compliance (such as in knowledge management). It is not clear to what extent, if at all, AI can or should solve structural issues that are rooted in non-compliance with policies and tools.

Busara is battling reputational concerns while also figuring out its own stance on AI. It is not clear if the best path is to be fully transparent about shadow AI use now, or try and streamline it first (with the known compliance challenges). This is exacerbated by the fact that the AI strategy implementation has an appropriately long timeline, but that shadow AI use is extensive.

All of the above issues are exacerbated by a lack of trust.

Box 3: Busara has a trust challenge

The consultations on AI unearthed a deeper and varied trust problem within the organization which this strategy with its commitment to transparency seeks to address. We see this trust problem as a broader pattern of how AI use without principles and guard rails is corroding relationships within organizations and across partners. People do not trust leadership to be able to draft and implement policies and then maintain compliance across the organization. Partly this is due to a lack of trust when it comes to the question whether leadership transparently discloses their own AI use (sometimes using AI checkers on org-wide communications in which no AI use has been declared). With the introduction of the sIfA tool, people do not trust that they might not be punished in their career progression—either by the organization or by their peers and donors outside Busara—if they disclose AI usage. People also do not trust their peers to transparently disclose AI use in teams, thus mistrusting work that has been shared. They do not trust their peers to accept Busara's AI policies and guard rails and expect them instead to use their own shadow AI to then present it as being produced within Busara's guard rails. This comprehensive lack of trust and active mistrust constitute a serious organizational challenge that Busara's leadership needs to openly address.

Change management measures

The AI strategy touches on every aspect of Busara, regardless of whether or not the outcome of strategic considerations is to use AI or not in a particular work flow. This means that implementation of the strategy is truly a shared journey, rather than a top-down implementation. It will be accompanied by excitement, anxiety, hopes and fears and will encounter enthusiastic application as well as resistance. All of these reactions are valuable, as they provide an accurate snap shot of what the arrival of AI means. Some of it is exacerbated since AI provides something of a moving target, highlighting the need to collectively own the value-driven approach of the strategy.
From setting up the strategy work, communication and involvement has been center stage (see 16.5 Appendix 5: Methods used in developing this AI strategy. Some of the communication channels listed below are already in place; others are new or could benefit from now becoming more focused on the specific tasks that the AI strategy requires.

Communications plan
Communications are building on the established consultative approach and focus on three message areas:

- Socializing the value-based AI goals
- Setting the learning and transparency culture
- Specific messages regarding implementation of the AI strategy

Various communication channels are used:

- Weekly org-update: for more detailed information that is relevant to everyone (weekly)
- Town hall meetings: For broad socialization of AI strategy knowledge (every three months)
- Jour Fixe: learning and exchange meet, with a write up appended to the org-email to broaden the culture of learning and transparency (monthly)

- AI sprints: AI-themed focus weeks to deepen learning, expand skills, exchange lessons and update on the AI-strategy KPIs (quarterly)
- Particular team meetings for asks or leadership requirements (SLT/ leadership meetings; portfolio meetings) (as per requirement)
- Executive team meeting to affirm exec level ownership of the strategy (updates fortnightly)
- Research festival Tara Mistari for more detailed hands-on workshop (annual)
- Co-creation workshops for specific tasks, governance developments such as guardrails (as per requirement, piloted in August 2026)

Monitoring and feedback mechanisms

To monitor change management, stay current with knowledge needs and collect feedback, each channel above periodically becomes a feedback-only session: no new information given, only feedback gathered. The Staff Pulse Survey and Leadership Survey add further feedback every three months. Together, the quantitative and qualitative data takes the pulse of acceptance and surfaces tension or disagreement with decisions taken during goal implementation.

An anonymous mailbox for whistleblowing or reporting misuse will also be set up, recognizing that AI use can make people feel vulnerable.

> *How will we support necessary discussions and changes brought about by AI, both in our external environment or for internal purposes?*

This strategy approaches the tasks ahead with ambition and skepticism, driven by values. Success first requires clear ownership by an AI culture lead—a transformation manager owning the philosophical, substantive and technical tasks the strategy poses. Who is working with adaptive and agile principles, featuring feedback loops and regular reflection built into the implementation plan to catch the needed adjustments, bottlenecks and flawed assumptions.

Further, a successful strategy requires that responsibility is spread wider through leaders within a project group having ownership of key goals, AI values ambassadors from each relevant team, and ultimately a wide spread of skills and consensus on AI use in the organization. However, it is important that in the setup of the project groups and AI values ambassadors, the team is made up of dissenting voices, ideally representing many of the archetypes offered in 9.4 Stakeholder archetypes.

The AI culture lead is supported by a knowledge worker and, where necessary, can draw on additional technical expertise either in-house or external. Table 4: Roles, task, candidates and time commitment for implementation of the AI strategy below lists designated candidates and time commitments. Training needs are indicative and will require confirmation in each case. However, it is important to note that currently no structured AI-relevant training exists.

Solutions and resources

Priority AI use cases

Use cases were identified through a mix of consultations, survey, executive team use case identification, and known pain points in the organization. The prioritized use cases below further lend themselves to become part of learning and growing AI skills in the organization, while also offering opportunities to address long-standing issues of compliance and consequences of non-compliance.

A number of use cases are explicitly out of scope, as per Busara's goals and five bold steps. These include: AI-supported primary field data collection (outside careful experimentation with AI-supported research methods (see Tier 2: mid-term use cases that need stronger foundations in improved data care or established AI governance principles), automating writing, checking outputs against a set of specified criteria to streamline a Busara voice, qualitative data interpretation with deep contextual knowledge, and any workflow requiring participant data that lacks specific AI-processing consent (which excludes all existing participant data up to now).

Use case 1: Critical thinking coaching
This is part of a broader training program honing critical thinking and AI skills.

Process: A coaching set up that becomes part of every individual's interaction with the platform (not a self-paced course that can easily be ticked off, but guard rails that challenge critical thinking steps during usage).

Use case 2: AI-related project work inventory
This regular scan allows maintaining the knowledge base on where AI is used in projects, preparing the information for the quarterly synthesis publication on AI and social science research with the Global Majority.

Process: Utilizing the principles of use cases 8 and 9 (surface but don't conclude validity of information; apply inclusion criteria but do not write prose), an agent scans and surfaces AI-related project work each quarter.

Use case 3: Transparency and AI governance
All Busara outputs mark the level of human/AI interaction through the sIfA tool (table and figure, or table only).

Process: All outputs going forward are automatically checked for a sIfA as part of a quality control process.

Use case 4: Meeting notes and action-item tracking
This is already an unsystematic practice across the organization and will be built out to be compliant with data protection standards and part of Busara's AI governance.

Process: After formulating our own and then aligning with partner's recording policies, meeting notes capability with an agent workflow is used to turn a transcript into a structured action-item summary with owners and dates, reviewed by the meeting organizer before distribution. The notes are used to build a project-relevant knowledge base, which gets archived after curating what outputs of a project are maintained within the knowledge database.

Use case 5: Project management
Using project management tools uploaded in the database, this streamlines project management patterns, allowing automated budget tracking, deliverable and milestone tracking.

Process: After curating a 'one RAP way' of project management into templates, each project is set up with a project space to automate tracking, linking meeting notes to action points and alerting to ethical or data privacy challenges. Closed projects are migrated into a closed shared drive after relevant documents for the knowledge database are identified.

Use case 6: Building a queryable institutional memory that enforces data protection and prioritizes AI on the input side

The knowledge management and data privacy problem has been a high priority for years, with an unmanageable backlog, compliance issues in saving folders, shadow data and a calculable cost of people needing to rebuild documents or instruments from scratch that they know do exist.

Process: Revoke all access to the shared drive given to non-Busara staff or former Busara staff. Sort the 1TB+ Drive using AI into outdated or duplicates (discard duplicates, file outdated or stale in a closed folder only accessible to management) using an agent's ability to create a simple inventory script. During this process, AI is also used to seek to identify root causes of flawed processes and clear documentation/strategy or policy gaps that exist within the organization. What is left as relevant is sorted into types of documents: org-policies, research evidence, project documents (such as budget templates), proposal support, with more detailed meta-data such as project, deliverable type, practice area, date, type of evidence—and possibly a rigor classification that could be applied to publications publicly going forward. This sets up the knowledge base that is organized for retrieval and reuse.

This subset is quality checked by relevant staff, including a peer-based review of which research evidence to include in the knowledge database. Curate that subset into a clean folder structure that becomes the curated knowledge base (up to 1,000 text-based files, 8M characters each). Connect a dynamic folder (with information likely to be regularly updated, such as publications, outputs or ongoing projects) via an agent-based folder sync (200 files/folder, 5 folders/agent, refreshed every 24h). Upload research evidence directly into the knowledge base to allow a stable and larger reference set of 800 documents), with an agent handling chunking and embedding automatically. Rebuild project folders only for open projects that can be shared with partner organizations. Disable sharing functions to only owners of folders.

Use case 7: Tagging/classification support for the outputs map
To tag Busara outputs (tools, playbooks, syntheses, reports) to practice areas and connect them to prior related work, gradually building a Busara evidence map and improved metadata.

Process: Using the infrastructure set up for use case 6, an agent with access to the outputs knowledge base suggests tags/practice-area classification and surfaces related prior outputs when a new document is added; a human (practice area lead or knowledge worker) confirms the tag.

Use case 8: AI-assisted public adoption alerts
This allows us to maintain a track record of Busara mentions beyond the more official Google Scholar profile.

Process: A scheduled workflow is configured to access public web data to produce a brief of where and how Busara is mentioned. The brief is reviewed by a human.

Use case 9: First-pass literature screening and literature review support
This seeks to improve on an established staff practice of using AI for a first pass at a literature review, which in practice is often done on unpaid tools without research mode and used to draft the actual review.

Process: An agent with access to relevant databases (SCOPUS, Web of Science, Open Science repositories) and a coaching set up that refines how the literature is queried scans abstracts and flags them against defined inclusion criteria. It does not write the prose of the review and also offers an entry for the sIfA.

Use case 10: Proposal/BD support
Using organization background knowledge, a weekly horizon scanning tool collects proposal opportunities and automates a triage process, sorting re-

quests for proposals into deadlines, scoring criteria, requirements, level of effort needed extracted into checklists and win probability. This replicates workflows that are currently done manually. The tool does not explicitly draft the full proposal, but in integration with the knowledge base offers standard proposal language, for example on AI use, and refers to relevant past proposals. In addition, under this use case sits the surfacing of standard contractual language on Busara's AI policy.[18]

Process: A workflow scans the horizon by querying relevant websites (such as DevEx or Development Aid, listservs but learning about more relevant sites as we go along), and lists relevant RFPs for a first human scan. With confirmed interest RFPs, the workflow ingests an RFP document and outputs a structured checklist (deadlines, eligibility, scoring weights) and enables a go/no-go decision-making process. While the tool can draw on relevant set pieces drawn from the knowledge database, writing the full proposal is explicitly out of scope to avoid hitting KPIs by simply producing more proposals with reputational risks attached. Substantive sections are always fully human written and marked as such.

Box 4: Tension: Can we not just improve our win rate by submitting many more AI-assisted proposals?

The organizational KPIs surface a number of core tensions relevant to the AI strategy. KPI 1 (total revenue, win rate) is made harder by unclear AI policies, where win rate might be dependent on the reputation of being AI-transparent, which might mean to not have used AI at all where clients are concerned that Busara is outsourcing thinking. The win rate, however, might be improved through AI-assisted proposal writing (simply because it allows submitting more proposals, not necessarily higher-quality ones), which cannot be seen as a clean improvement for this KPI. This tension remains even if revenue is in line with expectations.

18 This use case is drafted in collaboration with (Yalouris 2026)

Use-case suitability check

The use cases were chosen on the basis of a suitability check drawing on consultations and surveys, with more detail provided in the table below, in addition to value-based judgment that explains the inclusion of some weaker candidates. Yes = 1, Partial = 0.5, No = 0. The threshold for a strong candidate is ≥4/5.

Table 4. Use case suitability check

Use case	Deep pain?	Data available?	Process documentable?	Change measurable in 3–6 months? / Is AI really the best solution?	Score
1. Critical thinking coaching	Yes: 60.8% named role-relevant training as priority need; no structured training exists	Partial: the critical-thinking rubric is a Goal 1 baseline still to build	Partial: deliberately not a hands-off course; guardrails during use are hard to specify	No: literacy gain measured Q4 2028; design starts Q1 2027. But this is a fundamental issue for Busara. / Yes: AI literacy taught through the tool	**3.0**
2. AI project work inventory	Partial: real oversight gap, but serves a publication ambition	Yes	Yes: scan, inclusion criteria, surface without concluding	Yes: quarterly count / No: this is an AI ambition of importance rather than a use case	**3.5**
3. Transparency and AI governance	Yes: credibility challenge for global research	Yes: sIfA built and shipped	Yes: binary gate before release	Yes: Goal 2 targets Q1 2027 / Partial: AI tool and a compliance issue	**4.5**

Use case	Deep pain?	Data available?	Process documentable?	Change measurable in 3–6 months? / Is AI really the best solution?	Score
4. Meeting notes and action items	Yes: live shadow use with named risk sharing	Yes: transcripts generated at point of use	Yes	Yes: governed-meeting count / Yes: mature speech-to-text and summarization	**5.0**
5. Project management	Yes: both time and budget pains recorded	Partial: double bookkeeping; portfolio figures partly estimates	No: no consistent tracking process exists to document	Partial: adoption fast, outcome H1 2028 / Partial: AI here can help us in making this a focus area	**2.5**
6. Queryable institutional memory	Yes: 1TB+, 22 of 62 projects misnamed, calculable rebuild cost	Yes: abundant, if messy	Yes	Yes: files curated, 90% migration by Q1 2028 / Yes: no manual alternative at that scale	**5.0**
7. Outputs-map tagging	Partial: serves org-KPI 5; less acute than 5 or 6	Yes: published outputs, no consent issues	Yes: agent suggests, human confirms	Yes: % tagged, tag acceptance rate / Yes: classification is a core strength	**4.5**
8. Public adoption alerts	Partial: real gap beyond Google Scholar, moderate pain	Yes: public web	Yes: scheduled workflow, brief, human review	Yes: Goal 3 specifies adoption alerts show an upward trend / Yes: mature search-and-summarize	**4.5**

Use case	Deep pain?	Data available?	Process documentable?	Change measurable in 3–6 months? / Is AI really the best solution?	Score
9. Literature screening	Yes: already happening on unpaid tools without research mode	Yes: SCOPUS, WoS, open repositories (license position unstated)	Yes: query, scan, flag against criteria, no prose	Yes: throughput and precision against human check / Yes: screening at volume	**5.0**
10. Proposal/BD support	Yes: named major workload bottleneck	Partial: depends on UC6; Copper integration uneven; historical cost data lacking	Yes	Yes: Goal 6 targets 80% from automated scanning / Yes: RFP extraction and triage	**4.5**

Tier 2: Mid-term use cases that need stronger foundations in improved data care or established AI governance principles

Vector RAG

A Vector Retrieval Augmented Generation will be useful once a more structured insight into what knowledge exists and should be shared externally exists. This can include automated updating of the website search function, moving the process from a keyword/metadata search to an embedded semantic search.

Structured, queryable KPI/portfolio dashboard support

This could automate what is currently a messy and uneven process and thus first requires streamlining the process and its underlying data. In the future, a dashboard could automate queries on staff utilization, revenues, margins, and flag inconsistencies or risks.

Proposal scoping and level-of-effort estimation consistency

This is currently lacking clean data on how costly past engagements were, what role scope creep or overworking played, and how high the costs were of duplicated efforts. Improved data on this can in future enable better decision making on proposals and be included in the horizon scanning tool.

Careful experimentation with AI-supported research

As part of contributing to the knowledge base on what research AI can responsibly and reasonably support, anything involving frontier methods (such as synthetic-participant/agent research) is not part of a use case scope, as it does not include working with shared platforms and involves external academic partners, de-identified public datasets, and unresolved questions of consent and what constitutes ethical and contextually-relevant research questions. These will need to be resolved with an ongoing commitment through careful experimentation before considering any platform-based use case.

What human, technical and organizational resources are required?

AI skills audit

The most comprehensive oversight on what is happening on AI in the organization right now is this strategy—and it will have huge reporting gaps as participation in surveys and consultations was voluntary. This poses a high risk as AI use is disconnected, non-transparent, and does not contribute to crucial organizational and individual learning. The absence of AI skills also means that both staff and leadership underestimate what it means to set up the organization for AI. It is often expected that it is simply a matter of policies on how to use and not to use AI, whereas the reality is different: crucial processes need to be rethought with a view to deciding where AI can or should not be used. New regulations need to be taken into account. The executive team needs to be fully aware that they are the drivers of change management, realizing that the AI strategy touches on all aspects of the organization.

Primarily, the arrival of AI requires championing a culture of openness and learning, where time is invested into talking through processes, fears, risks, opportunities, trying out, innovating and establishing principles.

No structured AI training has yet happened. However, in consultations and in the survey, some individuals described themselves as above baseline skilled or senior developer level in AI; 43% reported feeling confident or very confident in their ethical use of AI. Skills are evidenced by having had some training or built agents; however, there is little evidence in the consultations that the high claims of ethical use are credible—only few ethical considerations such as diversity, exploitation, environmental concerns were brought up. AI disclosure, a cornerstone of ethical AI work, is absent and only slowly being introduced as part of the Futures initiative with the sIfA tool, which is Busara's own designed and shipped open-source software to create a Statement of Intellectual Fellowship and Accountability (see more at www.sifatool.org).[19]

Others reported using AI for specific skills within their jobs—such as quality checking for quantitative analysis, coding/debugging by, for example, running manually-written code past AI for review with the rule to not make changes, but give feedback only. Wider use cases were drafting/editing (71%), meeting notes (61%), brainstorming/self-checking (57% each), or summarizing (55%). Other types of knowledge work that came up in conversations were quality checks (for example on data management), revamping outdated proposal elements, identifying research gaps, and processing complex RFPs to extract deadlines, scoring criteria and requirements. With writing, people used AI as a writing coach to suggest structures (for example when writing proposals). Others use AI to digitize handwritten idea maps. Specific ideas on coaching include self-guided courses on relevant behavioral science concepts such as COM-B, and how to contextualize COM-B.[20]

19 (Schomerus 2026; Saleh and Schomerus 2026b)
20 (Saleh and Schomerus 2026a) (West and Michie 2021)

Tool familiarity is high (88% report ChatGPT familiarity), but displayed practices speak of shallow configurations. The lack of agentic or engineered workflows, context and guard rail setting or process clarity was evident. In the survey, 60.8% of respondents named 'training or examples relevant to my role' as a priority need. Self-identified training gaps focused on general AI knowledge, such as prompting or persistent instruction setting. However, learning on prompt engineering has been mainly a case of trial and error under time pressure, meaning progress was not prioritized.[21]

We currently have no knowledge on how AI comfort and adoption differs by tenure or seniority. There is speculation that senior staff are now using AI for work that they would have previously delegated to a junior staff member, stoking fears of deskilling, but there is no comprehensive insight into this.

There is currently no specialized advanced AI automation or engineered capability within the organization. The AI platform for testing what solutions for the organization could look like was set up by the author of the strategy, based on training through TÜV Austria, additional webinars, training on AI culture leadership through the Happy Company and extensive auto-didactics. While this is limited, it does at least follow the general AI strategy recommendation that for the implementation of an AI strategy, one champion sits on the executive level who is willing to expand knowledge on the topic and do the translation work for others.

However, there is a broader organizational need to understand the organization's work flows and what cultures underpin them (including people's fears and concerns) before introducing conversations on AI use within them—which is what the use cases in this strategy represent. Generally speaking, it is fair to say that process clarity is an Achilles heel for Busara, alongside record keeping and knowledge retrieval, highlighting the challenge of implementing AI within the organization's processes.

21 (Saleh and Schomerus 2026a)

Technical/platform resources

- Busara seeks a business-grade AI platform that can connect to existing knowledge sources (such as Google Drive and Slack) and can build simple agents and workflows without significant additional engineering capability. For year 1, such a platform offers a low-cost and ready-made configurable solution which allows a lower risk entry and supports learning AI skills across the organization, mindful of current technical capacities across the organization. This follows advice and best practice offered by competing organizations, such as IDInsight.
- Which model? Various platforms were reviewed, with preference given to one that offers highest GDPR compliance, including no data sharing for model training by default and hosting only on non-US servers. It is also model-agnostic but includes standard OpenAI, Google and Anthropic models but also open-source models such as Mistral. This means that Busara can choose which LLMs to use or to entirely exclude. Data sovereignty status can differ from model to model, so this is an important consideration.
- Knowledge base capacity curation/quality standard setting: Database caps mean that the knowledge management use case requires curation. This offers the opportunity to include human peer review (rating based on peer-judged quality) while also training and experimenting with AI-based curation when auditing the Google Drive backlog and using insights to feed into AI governance. This combined human/AI process will have to precede building the knowledge base to produce a Drive storage report before any work on the platform knowledge base.
- Data care: any document included in the platform knowledge base must be checked for personal data. No personal data can enter the knowledge base. This step will also help in further clarifying the scale of problematic data care at Busara and strengthen principles needed for AI governance.
- Workflow automation: for the triage in the horizon scanning model and the use case on adoption alerts, the workflow should always consist of

trigger/agent step/human review step. These become repetitive rather than ad hoc chat prompts, cutting down on inconsistent quality and duplicated effort.

Human resources

- Two new roles or redefined roles: AI culture lead (overseeing all implementation of the AI strategy, including more specifics like training, platform configuration, governance) and an AI-skilled knowledge worker (platform administration, platform configuration, data curation, metadata, tagging quality). The technical capability will need to be augmented with in-house capacity or external consultants where necessary. If AI use cases increase in scope and sophistication throughout the implementation of the strategy, this team resourcing will require revisiting.
- Each use case has a named internal owner. This means that responsibility for the human keeping AI in the loop is not centralized to an AI team, but assigned to relevant roles. For example, Practice Area Leads confirm tagging/classification outputs; the partnerships team owns the horizon scanning tool and reviews its outputs; Senior Directors in MRI and Directors in portfolios review all literature scans and contribute to improving the coaching and methods aspects of literature reviewing. Operations and Finance owns project management workflows, while compliance/IT/Software own all data care policies that control what can enter into the knowledge base. Crucially, each member of staff remains accountable for transparency (through sIfA) and for all text they produce as per sIfA and AI governance rules, as all AI-tools are restricted mainly to critiquing and surfacing standard text, not drafting text.
- Knowledge base curation: To make this a genuine Busara-owned peer-review driven process will require staff time to be allocated to the rating exercise.
- AI values ambassadors and goal owners: as per stakeholder analysis, a small number of staff per team (research, ops, BD) will require

dedicated time to co-create AI governance, co-build and troubleshoot all AI-facing configurations to support best practice peer learning. Goal owners will need to plan and resource implementation of their goals.
- All staff: all staff are required to participate in a to-be-designed AI Critical Thinker training, which uses different formats for different tiers.

Organizational resources
- Quick started AI governance: while developing detailed co-designed AI governance a quick guide as to what is absolutely not allowed to do with AI at Busara (and clearly outlined consequences) is a necessary first step. This will be an evolving document requiring updating and maintenance.
- Quick started AI skills: in order to get started on Tier 1 use cases, most staff will require some basic AI skills training (outlined below). These include things like prompting skills, judgement, guard rails but exclude deep engineering skills.
- Improve data and knowledge governance: this step adds to the workflow in each project (although in reality this workflow should already exist) to confirm the source of anything uploaded into the knowledge base, check that no personal data is included, and confirm a named human who reviews any agent output.

What skills does the team need?

- Prompting skills: any staff involved in configuring agents and workflows need to know how to write role descriptions, context, and instructions for specific outputs, as well as familiarity with established prompting approaches such as COSTAR. This is light initial training and can be provided by the AI culture lead.
- Curation and metadata judgment: this is not only a skills gap, but also a procedural gap within Busara since clear quality criteria either do not exist or have never been enforced. This skill set needs to be owned by the knowledge worker, who leads a training for the data curation

review team in order to curate the limited knowledge database to the highest standard.

- AI-output review literacy: everyone who works with AI in any way or capacity needs to know how to judge the quality of an output and recognize their own knowledge needs if they are unable to. This requires careful sequencing of AI use, only augmenting already familiar tasks and skill with AI, with further exploration only after basic skills are acquired.
- Governance and disclosure literacy: everyone using AI for their tasks needs to be familiar with how to use the sIfA and what disclosure practices it demands. This is a matter of culture change (transparency and learning are valued over efficiency) first and of skills second: they require absolute leadership commitment.
- Prompting discipline: the 'AI-in the loop' approach requires that agents need to be configured to explicitly coach and teach, critiquing and giving feedback rather than operating on the 'ask a question, give an answer' principle. This means that staff need to be skilled in prompting and configuring away from simple instructions to outsource work.
- Judgment on what should not be touched by AI/AI governance: staff working with personal data, using primary research skills for data analysis or analyzing literature running data-analysis or literature-review Agents need to have a principled and value-driven approach to deciding what does not become part of any AI interaction. This requires that leadership and AI values ambassadors role model and enforce these practices, that they become part of the AI Critical Thinker training, and that the platform itself flags uncleared practices (such as uploading identifiable or unconsented data or outsourcing cognitive skills).

Technologies and tools

Which AI technologies/tools are relevant?

Busara's AI strategy starts from three premises:

- We do not yet have agreed standards for our AI governance—part of the AI strategy is to co-develop these. However, we are driven by the need to sharpen our cognitive skills and use AI where it supports quality and innovation. This rules out a blanket approach to pursue as much AI automation as possible.
- AI skills across the organization are generally low. This means that for any AI use, skills training is necessary so that staff can understand not just a specific task, but where it sits in a broader landscape of AI tools.
- Processes and data care are messy and require a concerted effort to be addressed. Simply automating such processes and not addressing the lack of data care first would directly counter Goal 5's ambition to identify and fix flawed processes and Goal 2's ambition to establish responsible AI governance.

This means that for Busara's identified use cases, relevant AI technologies and tools are fairly narrow, basic, and specific. For Tier 2 use cases, the ambition moves towards a vector-based RAG. More details in Table 5 on the next page.

Table 5. What do different technologies do for Busara's use cases?

Technology category	What it does	Where it shows up in Busara's use cases
LLMs (chat/generation/reasoning)	**Search, classify, critique, summarize, answer queries in natural language**	Widely in Tier 1 use cases: note taking, coaching, quality checklists, horizon scanning
Retrieval-augmented generation (RAG) / in-memory embeddings	**Creates a numerical representation of a document, widening the options for an LLM to search and cite**	Mainly tier 2 when knowledge management exceeds tier 1 capacity needs
Agents (using an LLM for a defined task)	**Designed to enable an LLM to conduct a specific task with established guardrails and schedules**	Tier 1: Horizon scanning triage, feedback on research designs or outputs (not writing them), using literature searches to highlight collaborative potential across portfolios and projects
Workflows (using rules to run tasks; the rules themselves are not AI)	**Designed to follow an established process after a specified trigger (working with an agent might be one part of such process)**	Updating files in knowledge base, tagging/metadata on new files, setting reminders, scanning for adoption alerts or relevant RFPs, generating checklists
Web-search-and-summarize agents	**Distinct from using knowledge from Busara's knowledge database: searches the web (either entirely or pre-determined sites such as literature databases or funding sites) to offer an LLM-based synthesis**	Literature review research, adoption alerts, RFP horizon scanning
Meeting notetaking (speech-to-text and summaries)	**Converting audio into transcripts and structured notes, often with listed action items**	Use case on utilizing meetings notes for action items and follow up, within permission boundaries and data care

Technology category	What it does	Where it shows up in Busara's use cases
Consolidating structured data into a dashboard	**Allow to detect patterns and consolidate them across established KPIs. Can be AI supported, but is more likely to be non-AI, but supported by better data care and stronger process accountability due to AI strategy goals**	This is mainly a problem of data gaps and record keeping that requires more diligence on either, rather than an AI solution. However, part of Goal 5 processes will be to determine if some processes can improve by establishing and maintaining robotic process automation (RPA) to avoid entire systems changes

Some technologies are explicitly out of scope for Busara until further use cases are clarified, AI governance is determined and AI skills are improved. This applies to agent architecture with autonomous unsupervised memory management. Currently, image generation has been in use for a while and requires transparency, but not necessarily platform integration.

Hosting, licensing, adaptability and architecture patterns

Busara takes a cautious and informed path towards learning how best to use AI for our workflows. Through the skill-building phase, which includes all Tier 1 use cases, the default is to buy and configure software as a service (SaaS). That means fewer customization options at first, matching current skill levels (including on security, compliance and LLM licensing) and the priority of fixing flawed processes. Decisions such as building a vector database are deferred to experience and learning in Year 1.

Our needs point to a cloud-based, integrated, GDPR-compliant, EU-hosted and model-agnostic platform, which is The platform's default setup, and that of any LLM it routes to, with exceptions on server location clearly noted. Un-

paid versions of any OpenAI or Anthropic product are not permissible for any task relating to Busara internal knowledge (which however, is likely regularly violated). Self-hosting only becomes economically viable, per hardware calculation, at much higher token use and compute needs than Tier 1 requires.

Year 1 focuses on configuration skills and learning, matching the needs of all Tier 1 use cases. If the ceiling of The platform's knowledge base proves too low for Busara's purposes, a growth path towards more customized solutions is designed. Switching costs are low: moving means rebuilding rather than losing data, since Busara continues to own all of it.

Through the skill-building and learning phase, Busara's default position on agentic architecture is that all agent work is explicit, checked by humans at defined steps, and always legible to humans with basic AI skills (which follows a LangGraph pattern, not a MemGPT pattern in which agents decide what is salient and what can be discarded as irrelevant).

EU hosting may require revisiting once Kenya's AI Bill passes: cl.36(2)(e) names future rules on 'data localization', and Pillar 2 of the Kenya AI Strategy 2025–2030 names 'data residency requirements'. Neither has passed into law, so EU hosting is currently a good choice. But Kenya's AI legislation, and that of other territories relevant to Busara, requires monitoring.[22]

Technologies and tools market scan

None of our competitor organizations listed earlier have reported building a custom LLM infrastructure, but rather choose similar pathways (configuration) to address similar pain points (knowledge management). IDinsight explicitly recommends committing to one platform, rather than patchworking many different solutions, which is the approach we are mimicking.

Retrieval/Knowledge-base platform options are outlined in Table 6 next.

22 (The Republic of Kenya 2026)

Table 6. Retrieval/Knowledge-base platform options

Option	What it is	Why considered / rejected
Selected platform	**Chunking, embedding, retrieval, up to 1,000 files, no infrastructure**	Requires no infrastructure building, open source LLM is included, offers enough scale for our current needs, avoids patchwork solutions
Self-managed vector databases (Pinecone, Qdrant, Milvus, Azure AI Search, AWS Kendra, Vertex AI Vector Search)	**Requires Busara to run the embedding pipeline and index; could be integrated into The platform, but spreads the solutions across tools**	Too soon for our early phase, as it requires pre-empting AI governance decisions on which embedding model to choose, requires AI skills to build and maintain an ingestion pipeline which requires a designated technical team. No advantage until we hit the The platform knowledge base ceiling
Third-party drive-audit tools (CloudM, Netwrix, Alphamoon)	**Governance dashboards for duplicate or stale-file cleanup at scale**	While an attractive option for our huge knowledge management backlog, it assumes that most of our backlog is relevant, which it likely is not. Since it does not support retrieval, The platform functionality will suffice for a broad identification of duplicates and stale files, while Google Drive's built-in storage manager plus a simple inventory covers current needs

Meeting-notetakers: Widely in use in Busara are Zoom recording, Otter.ai and Read.ai, with Motion sometimes used. Otter.ai and Read.ai both carry risks, of which some staff are aware: Otter retains prompts for 30 days and shares these with AI and analytics providers as well as law enforcement agencies. Some staff have already taken the step to no longer use Otter. Read.ai is not universally trusted. An AI governance policy will likely focus on integrated solutions—for example using Zoom's and The platform's inbuilt compatibility, where The platform pulls the zoom recording to extract insights and action items and sends these by Slack (if integrated).

Workflow-engine comparators: These are not yet relevant for Busara's use cases—none of the Tier 1 use cases requires orchestration across systems.

Solution levels and caveats

Mapped against solution levels, most Tier 1 use cases sit in the 'buy software as a service' tier: costs are The platform licensing plus staff time for configuration such as knowledge base assessment. Surcharges for agentic work involving judgement calls, and low-code requirements such as the horizon scanning tool and the adoption alerts, are accommodated within this level. API integration or custom builds are a question for later, likely at the end of the period this strategy covers.

Two Tier 1 cases that are not software as a service are the sIfA tool and Busara's conference tool.

Any solution for the identified use cases will be finalized only after a trial period of up to six weeks, testing that the match of use case and solution delivers a minimum viable product (MVP). This includes testing actual token costs for the work we expect, so that realistic budget considerations can be updated against monitored usage.

The trial period also serves as a quality test case. A red team needs to check what the triage of the knowledge management backlog produces, whether the curation process yields reliable outputs, and what hallucination controls may need to be put in place.



Data strategy

Digital and online safeguarding of data (as stipulated in Busara's Safeguarding Policy) need to be front of mind. This means that:

1. ***Where digital tools or remote engagement are used, Busara shall:***
   - Use secure communication platforms and systems;
   - Protect participant identity and digital privacy;
   - Minimize risks of unauthorized access or exposure;
   - Prevent online harassment or exploitation; and
   - Apply appropriate cybersecurity and access control measures.

2. ***Personnel shall:***
   - Use approved communication tools;
   - Protect passwords and devices;
   - Avoid sharing sensitive information through unsecured platforms; and
   - Report digital safeguarding concerns immediately.[23]

What data sources are available?

- Published Busara outputs: these are hosted on the Busara website and include Groundworks, Aha! Moments, playbooks, tools, project reports in old formats and links to some academic articles. These link to org-KPI5/Goal 4 and can be included into an AI-supported knowledge management platform without further criteria.
- Internal knowledge: Busara's main repository for internal but non-confidential documents is a shared Drive.
- Slack as a data source: Slack is often used as a repository for notes and documents.
- AI note takers as a data source: AI notetakers notes are not stored in any systematic way.

23 (Muange (Manager Project Operations and Busara ISERC) May 2026, 22-23)

- Strategy documents: strategy documents (from org-strategy to portfolio or departmental strategies) are not readily accessible or non-existent.
- Raw research participant data: This confidential data (governed by Kenyan data privacy) is stored on a separate raw data shared drive. Any inclusion even in an internal AI-supported knowledge management system would require obtaining updated consent.
- Contact management, proposal submissions, proposal win rate: These are housed on Copper, with various levels of integration into staff email.
- Shadow data: There is an unknown volume of relevant project data, contact management, notes held on private drives that are not part of any shared system.
- External public sources: These include academic literature, partner publications, strategy documents. These are relevant for knowledge searches on, for example, the state of the field, but are not part of an internal knowledge management strategy.

Data requirements by use case

To follow best practice on data minimization, preparing relevant knowledge for AI processing explicitly excludes data that is not necessary for an identified use case.

Tier 1 use cases

One data exclusion applies to all Tier 1 use cases: no raw participant data or non-anonymized personal data is ever included. Sensitive financial data is included only under restrictions, such as being limited to certain accounts.

Use case 1: Critical thinking coaching

Data requirements and sources: This requires iterative configuration of the platform, with refined prompts, queries, guard rails and coaching processes on how staff interact with AI. While no additional data is required beyond what exists from other use cases, this might lend itself to some data extraction from platform usage logs to assess anonymized and aggregated use patterns of how staff work to refine the coaching over time.

Exclusions and human judgement: This use case does not explicitly require human judgement, but is a human judgement training tool (with coaches making some meta-decision on course design). What this use case proposes constitutes an explicit and to be decided culture shift for Busara, where staff is aware that their platform interactions might be assessed (anonymized) to refine coaching and guard rails. This is a transparency requirement, but might present to the organization as a privacy issue, where people might feel surveilled (based on lack of trust and the experience of people feeling unsure if their anonymous staff surveys or questions are indeed anonymous).

Use case 2: AI-related project work inventory
Data requirements and sources: Metadata (sIfA and project tracking) drawn from all internal project records tagged as involving AI in the research design (not in the processes covered by other use cases).

Exclusions and human judgement: Humans assess which flagged AI-related project work is relevant for the quarterly synthesis publication (which is fully written by a human).

Use case 3: Transparency and AI governance
Data requirements and sources: sIfA-based metadata for each output going forward (how much AI was used 0-2, which tool, for what task and with what reasoning). This is generated as a live document for all collaborative outputs, with staff using sIfA to self-report, augmented by AI notes when feasible.

Exclusions and human judgement: Humans are accountable for providing an accurate sIfA. sIfA is deliberately designed to make reporting as easy as possible and can be further refined as we learn.

Use case 4: Meeting notes and action-item tracking
Data requirements and sources: Draws on transcripts, attendee list, diary listed agenda, prior action items from internal meetings as well as meetings involving external partners.

Exclusions and human judgement: Use case 4 requires a clearer consent policy, including aligning with external partners' recording policies, and the end of shadow recording (such as unseen Granola). If the purpose of a meeting is to gather data (such as phone interviews), that data is treated as raw data that does not form part of the knowledge database. Humans judge the quality of the transcript and accuracy of the action points. AI is used to compile project-relevant summaries that are judged by humans for their archive value; raw transcripts are deleted after a specified time.

Use case 5: Project management

Data requirements and sources: Management documents that include project budgets, burn rates, deliverables, schedules, linked meeting notes and action points from use case 4, a to-be-developed 'one RAP way' template. Requires streamlining existing documents that are now different between projects, project tracking tools and ongoing input from the Finance team and operations.

Exclusions and human judgement: Sensitive financial data is not accessible org-wide by default but to specific authorized staff accounts. Upon project closure, projects are audited for relevant material to be included in the knowledge base. Humans decide on that long-term relevance, budget and scope and escalate any ethical, compliance or data privacy concerns. AI scans for budget burn rates, milestones, missed action items and sends alerts.

Use case 6: Building a queryable institutional memory that enforces data protection and prioritizes AI on the input side

Data requirements and sources: Up to 1,000 text-based files and up to 800 documents in the knowledge base and dynamic folder sync curated from a subset of Google Drive content and other actively managed project folders/ Slack: published/unpublished Busara outputs, organizational management data, policies, project and proposal documents (budget templates, proposal templates, set pieces). All curated documents require adding metadata, such as project, deliverable type, practice area, date, type of evidence (and additional metadata as emerges through consultations).

Exclusions and human judgement: Curating this data entails a major organizational infrastructure project, as the first pass at the Google Drive is to triage outdated, duplicated irrelevant documents into an archive-only drive that is not readily accessible for the organization. The knowledge base is not accessible to external partners or former Busara staff. Quality check and relevance for final inclusion in the limited database comes from a peer-review judgement call.

Use case 7: Tagging/classification support for the outputs map
Data requirements and sources: The knowledge base built for use case 6 includes metadata for practice areas that allow surfacing of existing practice-area relevant material.

Exclusions and human judgement: Humans confirm whether practice area relevance is correct and whether an output can be considered offering strong or weak evidence.

Use case 8: AI-assisted public adoption alerts
Data requirements and sources: Public web data such as external publications, citations, mentions of Busara's work.

Exclusions and human judgement: This excludes all data internal to Busara. Before anything counts as evidence of adoption, humans validate it to assert in alignment with project team and partner that a mention represents genuine adoption.

Use case 9: First-pass literature screening and lit-review support
Data requirements and sources: This requires access to academic databases (SCOPUS, Web of Science, Open Science repositories) and all listed abstracts (ideally full access). Some of the databases are open while others require subscription.

Exclusions and human judgement: Some databases restrict automated querying. Search is conducted with human-defined inclusion/exclusion criteria

against which AI only flags. The process feeds into the sIfA as an automated step or matter of principle. Humans structure and write the literature review.

Use case 10: Proposal/BD support

Data requirements and sources: Horizon scanning of web-published or emailed RFPs (DevEx, Development Aid, UN/World Bank procurement); past proposal, capability statements, relevant policies, non-confidential rate cards; set pieces drawn from the knowledge base.

Exclusions and human judgement: No data or outputs marked as confidential or proprietary. Substantial parts of the proposal (research design, situating in the literature, relevance to RFP) are human written. Humans decide on go/no go for proposals, assess the suggested win probability. AI only extracts facts like deadlines, eligibility, scoring weights, suggests probable matches, creates checklists and reviews outputs for feedback, not prose writing.

Tier 2 use cases future data needs

- Vector RAG: mature data from use cases 6 and 7
- Structured, queryable KPI/portfolio dashboard support: clean and reliable data on staff utilization, revenue, margins
- Proposal scoping and level-of-effort estimation consistency: clean and reliable data on historical cost, scope creep, cost of duplications
- Careful experimentation with AI-supported research: anonymized and publicly shared datasets under open science terms with explicit AI-related consent with strong safeguards and public reflections on the value of such research; ethical emphasis on making this research relevant and meaningful to participants.

Improvement needs for data quality

This data strategy is only valuable with maintenance of all data and rigorous quality checks. Data quality and accessibility is one of Busara's pain points—a crucial existential issue for a research organization—and persists despite

broad adherence to data policies covering all jurisdictions in which Busara works, such as the Kenyan Data Protection Act (2019). This strategy might be a first step towards ISO certifications. Quality checks must therefore include regular audits owned by the knowledge worker, provenance checks, metadata and classification, defined rules on what to keep and discard, and documentation of known biases and of data with reproduction challenges.[24],[25]

Streamlining Busara's data practices also requires merging informal data practices, such as using Slack for document storage, making relevant documents easily accessible.

24 (The Republic of Kenya 2019)
25 Drawing on (ISO 2023)

Security and ethics

Anything we do in developing our approach to AI within the five goals needs to be compliant with existing Busara policies on ethics, safeguarding and data privacy, as well as being able to be flexibly adjusted to respond to evolving legislation. Busara's recent NACOSTI accreditation as an IRB-granting body has highlighted that most of our external facing practices pass muster. Training on the Data Protection Policy is currently conducted by the compliance team (79 staff trained, 70 remaining), however, the training has struggled with low attendance and embedding and maintaining any policy compliance is a known challenge.[26]

However, additional demands that emerge from AI require further tightening or discussion on additional practices related to security and ethics.

Busara's AI governance structure

Currently, AI is barely mentioned in any policies: there is no AI-specific policy or data policy focused on AI, however, existing data policies touch on issues of AI. This is a crucial credibility gap, as Busara is not able to credibly report to partners how AI is used. Some clients impose AI rules, but there are no verification procedures in place to assure that these have been adhered to, which poses a major reputational risk. While no data handling incident related to AI use has been reported, it is perceived as a live risk by staff (and is likely to have happened in shadow AI use). Any future work on AI governance needs to be embedded within existing and expanding research governance (through MRI) and compliance (through operations).

Some fledgling accountability mechanisms exist, namely the shipped sIfA tool that offers a user-friendly and enforceable starting point towards a culture of transparency. In addition, data protection has been recently strengthened after a period of understaffing and is now overseen by the Director of Operations.

26 As reported by the Manager, Project Operations and Busara ISERC, July 2026

Data policies

Busara's current legal baseline is Kenya's Data Protection Act (2019). Busara's data protection policy was built to cover the requirements of every jurisdiction in which Busara operates. The bottom line is: no participant data can ever be processed by AI. The only change to this might come when two things change (meaning, Busara decides it wants to change these):

- Changed consent forms
- Consent to allow data to cross borders (which may not align with future Kenyan AI Act provisions).

Specific clauses that relate to AI in Busara data protection policy stipulate that every data subject has the right to not be 'subject to automated decision making: e.g. Individuals have the right not to be subject to decisions made solely by AI systems such as being rejected for a research study or profiled for targeted outreach without human involvement or the opportunity to challenge the outcome'.[27] Further, Busara's policy stipulates that a project-based data protection impact assessment assess if any 'data processing activity... is likely to result in a high risk to the rights and freedoms of individuals. This includes, but is not limited to, projects involving large-scale processing of sensitive data, use of new or emerging technologies (such as AI-based profiling tools), or systematic monitoring of research participants'.[28]

This also means that Busara will not allow its collected data to be used to train AI—certainly not without new and changed consent, but likely not at all as a matter of principle. This does not contradict Busara's commitment to open science data sharing, which is for scientific rather than commercial benefit.

27 (Muange (Manager Project Operations and Busara ISERC) September 2025, 13)
28 (Muange (Manager Project Operations and Busara ISERC) September 2025, 26)

Current consent forms, such as the ISERC Adult Informed Consent Form contain no AI-specific language and offer a broad clause on future use of de-identified data, which in legal interpretation clearly does not include specific consent to future AI-related data use. A practical consequence for Busara is that none of its existing datasets can be used for any AI-assisted work (including storing and training models) without requiring new consent. Since this is not practical or realistic, existing data sets simply cannot touch anything AI related. Whether new consent language should address this is a matter of discussion within the AI governance goal. In addition, most existing agreements prohibit sharing data with AI companies (which was a specific discussion point during Busara ISERC's NACOSTI accreditation process).[29]

Cross-border transfer of data is a concern, as most AI providers process data outside Kenya, which requires new safeguards. However, this insight might also uncover a loop hole in Busara's data handling, with data currently stored on servers outside Kenya.

None of the tier 1 use cases involve research data processing. However, future evolving/Tier 2 use cases will require stricter data processing to establish if personal data is involved, whether it has been anonymized, whether valid consent for the type of processing was given, and whether data will likely cross borders for analysis and storage.

Existing policies, gaps and fragmentation on security and ethics

While org-wide, detailed and specific AI guidance is lacking, Busara's Research Ethics Policy specifies 'Ethical Considerations in Emerging Technologies', which point forward to the need to develop guidelines:

The use of digital tools, artificial intelligence (AI), biometric sensors, and other emerging technologies in research introduces unique ethical considerations. When using such technologies, researchers must:

29 Information provided by the ethics and compliance team, July 2026

- Ensure full transparency with participants regarding how their data will be collected, processed, and used.
- Actively work to identify and mitigate algorithmic bias in AI-assisted analysis or decision-making.
- Minimize risks linked to digital surveillance, data aggregation, or re-identification.
- Seek IRB approval for any new, experimental, or high-risk technology-based research methods.
- Ensure participants are informed about and able to comprehend the nature of digital data collection.
- Implement additional layers of data security appropriate to the sensitivity of digitally collected data.
- Comply with relevant technology-specific regulations and guidelines as they evolve.[30]

Other fragmented non-org wide guidance exists, with several departments (BRACE, Social Cohesion) having developed their own rules for how to use AI (see 16.2 Appendix 2: PROTECT Code of Busara's SCP), or having had brief conversations about disclosure guidance (e.g., MRI and FIVE). Some client-imposed rules when it comes to AI use exist, with some named clients having strict no AI policies or other handling restrictions for their commissioned projects; however, these will be applied ad hoc on specific projects.

Most policies dealing with data focus on empirical research data, rather than offering a comprehensive data strategy. A data protection policy covers research data management 'relevant to all personal data processing/handling activities'.[31] Guidance for knowledge management, how and where to store documents and other issues of data handling is noted in Busara's IT Operations & Systems Policy; however, compliance and enforcement of these guidelines has not been a success.[32]

30 (Muange March 2026, 12-13)
31 (Muange (Manager Project Operations and Busara ISERC) September 2025, 11)
32 (Omedo (People Officer IT Systems and Support) November 2024); This draws on interviews with the compliance and ethics team

Policies on how research data cannot be used in relation to AI are clearest, due to the legacy of data protection policies prior to AI introduction. For internal use of data as it relates to AI (e.g., an anonymized, queryable repository with AI-assisted analysis), current guidance is the following:

- Data protection law doesn't prohibit the use of AI; it regulates how personal data is processed, so the first question is always whether the extract contains personal data.
- If it does, it needs to be anonymized/de-identified before going into an AI platform.
- Processing requires participant consent.
- Cross-border transfer (which most AI providers engage in) requires either proven adequate safeguards or explicit participant consent; this is the area that requires the most caution.
- Higher-risk processing may require a DPIA.
- Sensitive personal data carries additional cross-border restrictions.[33]

Busara's IT Operations & Systems Policy stipulates drive storage guidelines and data handling has a cyber security roadmap document, but no broader IT or data strategies to outline how compliance needs, organizational ethics, training and IT infrastructure intersect to build a comprehensive approach to Busara's knowledge management and the skills its upkeep and compliance requires.

This strategic and policy fragmentation is not surprising, considering that Busara's decentralized decision-making, with directors having discretion over team practices, creates limited central oversight, and limited compliance with organizational policies.

33 (Muange (Manager Project Operations and Busara ISERC) September 2025)

Box 5: Tensions in strategy alignment and prioritization

The org strategy does not address the tension that resource constraints (the need to generate revenue, shrinking budgets coupled with stronger client demands, limited allocated time for business development) tend to routinely override strategic considerations, and that internal processes are often disconnected from or misaligned from broader strategic considerations.

To what extent is US or EU AI legislation relevant to Busara?

While Busara works a lot with US funders, the US regulatory landscape on AI is inconsistent, with no federal act regulating state level laws. It is thus safe to assume that even US funders will be reassured if Busara orients itself on EU/Kenyan regulatory standards.

While Busara is mainly a user of AI systems and far from using high-risk AI tools, various parts require paying attention to. Busara could inadvertently become a provider if it builds AI-driven tools. However, any Busara-built AI tool (such as sIfA) is unlikely to fall into any high-risk category, but awareness of the possibility is necessary. Should Busara ever veer near high-risk frontier tools, Art 6-49 stipulate technical and organizational requirements on quality, data and risk management, testing, monitoring, documentation and accountability responsibilities. Incident reporting duties (within 15 days of noticing an incident, 2 days for wide-ranging/critical-infrastructure incidents, 10 days in case of death) needs to be mapped against Busara's existing IT incidence reporting practices, which are currently unclear and against necessary reporting procedures as stipulated by Kenyan law or other relevant jurisdictions. Failure to report incidents can come with serious sanctions.

Of relevance in the EU AI act is article 5, which lists practices against which a behavioral research organization needs to actively screen: manipulation of behavior, exploitation of vulnerable persons, social scoring, and specific biometric identification uses (e.g., racial profiling), alongside emotion inference in the workplace or behavioral targeting that could be classified as seeking to manipulate behavior. Article 50 stipulates duties on transparency and how these apply to emotion recognition, biometrics and generated content.[34]

Kenya's regulatory environment

Kenya's 2026 AI Bill (not yet passed into law) is a functional equivalent of the EU AI Act and mirrors it in most of the fundamental questions, such as risk classification, human rights impact assessment, cybersecurity requirements, ethics guidelines on bias, data privacy, environmental impact assessment. Kenya's draft also includes the stipulation that AI 'enhance rather than replace human capabilities', with mandatory human oversight for decisions affecting rights, safety, or wellbeing (p.32), which acts as a strong regulatory anchor for Busara's AI governance.

As regulations change or are confirmed, Busara needs to be in a position to flexibly adjust. This is made possible by this strategy's approach to be value-driven, establishing live documents over a rules-and-policing approach and needing to establish focused work groups with the relevant actors.

Risks

Despite fighting bias in social science research being Busara's *raison d'etre*, the risk of producing biased research when using AI tools was not named as a major concern. Countering this requires concrete steps, such as proposed by *Masakhane* through community-led reviews to catch bias and marginalization. Internally, part of Goal 4 is to establish a designated review step to check for bias. When choosing which LLM to use, Busara needs to test in a trial phase which models fail or deliver and how.

34 (European Parliament and Council 2024)

Working with AI poses another ethical dilemma for Busara: we are working across Africa, Asia, and Latin America while using technology (including AI) largely designed by and for the Global North but often data-fed and checked by the Global South. We cannot deny that this reinforces power imbalances and knowledge hierarchies.[35] One starting point to rectify this would be to prioritize the Busara research agenda on AI and the Global Majority, link into the relevant networks seeking to counter this, such as BenevAI, and put in effort for Busara's thought leadership visibility.

- The risk of deskilling is addressed through Goal 1 with training on AI literacy and critical thinking, as well as the guard rails put in place around writing and data work.
- Reputational risk is addressed through Goal 2 on AI governance, including transparency, and Goal 4 on publications to in future not publish work that is largely AI written.[36]
- Flattening of voice through AI use is addressed in the ambitions of Goal 4.
- The emerging problem of distrust among staff is at the center of this AI strategy and the full commitment to a culture of transparency and learning.
- Environmental cost is real and often discarded. In this strategy, it is covered by conducting environmental audits for AI use (for starters in use case 6 when keeping a record of token use during the trial phase) and with plans to publish an environmental footprint. In establishing guard rails, we will also reassert that not every task requires AI such as spell-checking.
- Risk of extractive research: research with AI is at the same risk of being extractive as other research. This is particularly prominent for the kind of work Busara does with vulnerable and often marginalized populations. It is exacerbated by the often-huge knowledge gap that participants have on what AI is, which likely makes informed consent impossible.

35 (Hao 2026)
36 How to manage specifically reputational risks needs to be guided by Busara's Risk Management Policy (Muange (Manager Project Operations and Busara ISERC) June 2026)

This stresses the need for extra careful assessment procedures on any research that touches on AI, with stronger questions asked about relevance for affected populations.

Responsible AI use in Busara

As one of the first deliverables, an evolving staff practice code will be built, which reiterates the five bold steps as core values and can take inspiration from SCP's PROTECT code (which the team developed and refined with AI assistance) and also establishes clear AI-free zones (full PROTECT Code in the Appendix).

P — *Proportionate use*
R — *Review it, always*
O — *Own the decision*
T — *Transparency*
E — *Expertise*
C — *Confidentiality*
T — *Trust but verify*



Revenue potential and cost efficiency

The cost of time lost to searching for information

The staff survey on time spent searching for information solicited a very small response (n=6); however, it confirmed what is a common narrative and experience in Busara: that finding information is one of the single-most frustrating and self-sabotaging experiences that is systemic across the organization (and not a matter of training). What the small survey indicates, however, is that there is variety in search times. The best-case scenario ranges from ten to 60 minutes (2 answers). The worst-case scenario is severe, costing people half a day. One respondent even wrote 'there's a document I've been looking for for about 3 months and still can't find.' Two respondents highlighted that they redid work that they knew to exist—one said 'very often' while the other offered concrete examples of redoing work after giving up searching. The cost to the organization of searching and duplicating effort scales with seniority.

Revenue potential

The AI strategy currently offers no direct path to new revenue. The possibilities are indirect, or recouped opportunity costs: thought leadership backed by a research agenda on AI might develop new partnerships, and an AI-literate critical thinker training approach might become monetizable, though it would need piloting as an offering. A figure for revenue potential therefore cannot be established.

However, there are indirect links to revenue potential that come from capacity reallocation, tied to org KPI 1 on win rate.

- BD efficiency: If AI-supported knowledge retrieval, RFP horizon scanning, proposal scoping and better set pieces free up time currently occupied inefficiently and haphazardly, freed up capacity can go toward more targeted or better-scoped proposals (not just more proposals).

- Time for relationship building: Freed up BD time could be spent developing relationships that are crucial to the emphasis on long-term partnerships, feeding into higher renewal rates and a stronger share of multi-year partnerships.

1.1 Scenarios

Avoiding losses by avoiding budget overruns

Table 7 below outlines some possible scenarios.

Table 7. Qualitative scenarios for avoiding losses from budget overrun

Case	Assumption	Suggested possible effect
Worst case	**Tracking tool adoption stalls; Hub automation does not materialize; overrun pattern continues at current 24% off-budget rate**	Loss is materialized
Expected case	**Excel tool creates some improvement in practice but will likely not comprehensively address frustration; rate of projects on budget stabilizes or slightly recovers back to 80%**	Some loss is avoided
Best case	**Tracking tool automation sticks (and is made more reliable through AI use) across all portfolios, closing the difference in practices gap**	80% or more of projects close on budget with overrun losses a rare exception

1.1.1 Knowledge management scenario

The calculation scenario below is based on rough estimates due to the low response number in the survey.

Table 8. Annual value recovered through time saved on information searching

Case	Share of search/rework time recovered
Worst case	15%
Lower-middle case	25%
Break-even	~33.8%
Expected case	40%
Best case	50%

Assumed benefits are:

- BD proposal efficiency leads to higher quality proposal with a higher win rate (feeding into org KPI 1 on revenue).
- Knowledge management system encourages reusing existing information, lowering duplication of effort (feeding into org KPI 2 on margins).

Even below a cost-efficiency threshold, addressing the known frustrations about knowledge management is valuable. What the effect of a platform looks like will need validating through a follow-up survey and an analysis of usage patterns. A functioning knowledge base on the platform pays for itself if it saves each of the 140 staff members just 15 hours of search time (roughly two working days) per year at the Analyst internal rate. However, it is important to note that this does not reflect the staff time necessary to set up the knowledge base.

Implementation plan

This implementation plan leans into the phasing of Busara's org strategy, where Year 1 builds infrastructure. This corresponds roughly to the Tier 1 use cases, spread over 18 months, which allow assessment of whether they bring measurable value within three to six months. If they do not, the sequencing offers a moment to pause and reassess. An overview implementation plan based on AI strategy KPIs is provided in 16.7 Appendix 7: AI strategy rollout swimming lanes

Project launches

Six evidence-generating and culture setting projects are necessary because piloting Tier 1 use cases requires evidence about how AI use affects Busara's voice, rigor, skills, and internal consensus on AI governance as the foundation for Goals 1 and 2. Without this step, Tier 1 pilots risk scaling a tool without knowing whether it protects or erodes what makes Busara distinct.

Table 9. Evidence-generating projects

Project	Goal	Timeline	Contributes to which use case or AI strategy goal?
AI voice checklist	Ensure that any AI-in-the-loop work does not flatten Busara's voices, depth and rigor. Define what that means through staff interviews, comparative analysis of AI-assisted and human-only reports, and a piloted brief on how AI could enhance or erode Busara's voices	7 weeks	**Goals 1 and 4**

Project	Goal	Timeline	Contributes to which use case or AI strategy goal?
Do this/ don't do that	Draft initial guidelines for AI use, data protection, transparency, cognitive guardrails for AI use, build on team guidance that already exists in Busara	8 weeks	**Goal 2**
10,000 hours skill protection plan	Map which mundane practices build foundational skills (writing, coding, pattern-spotting) and design AI configurations to augment (not replace) them; map what trains critical thinking skills and develop a syllabus	9 weeks	**Use cases 1 and 2; Goals 1 and 2**
AI ethics and practices debates	Internal debates surfacing staff practices and disagreement on AI use to foster a culture of learning	8 weeks	**Goals 1 and 2**
AI research agenda	Build AI-related project inventory, gather material for quarterly synthesis publication, develop format for the publication, develop research agenda	4 weeks	**Goals 1 and 4, use case 2**
Quality criteria for literature reviewing	Agree on quality control and coaching steps for literature reviewing	8 weeks	**Goal 4, use case 9**

Two infrastructure projects need to launch in parallel as prerequisites for use cases:

- Drive audit and curation project (part of Use case 6): inventory the knowledge management backlog using the built-in storage manager or a simple Apps Script report, decide relevance filters (last-modified date, active owner, file type, duplicate detection), and curate a subset into a clean folder structure. Its scale and level of effort make it a named and resourced foundational project in its own right, setting the quality standards that Use case 6 and org-KPI 5 depend on.

- Filing convention and metadata design project (linked to Use case 6): define tagging and classification conventions (project, deliverable type, practice area, date) on which org-KPI 5's outputs map and use cases 6 and 7 depend.

Processes that need to be adapted, socialized or created

The table below outlines the procedural status quo and what needs to be added.

Table 10. What processes exist and what is needed

Process	Status today	What's needed
Project tracking	**No consistent 'one RAP way' tracking process exists**	A minimal, consistent project status/budget-burn tracking routine, coordinated with the delivery function's Year 1 buildout of non-AI supported financial tracking and quality reviews
AI disclosure and transparency	**Launched, needs practicing**	Enforcement on all public-facing outputs, socialization within the org for all processes to build trust and a culture of transparency
RFP response to AI-related clauses	**No standard text exists to respond to RFP conditions on AI use (explicitly requesting, requesting disclosure, prohibiting)**	A short internal process to sign off on texts that is rooted in Busara's AI governance and responds to the three RFP scenarios, feeding into the horizon scanning tool (use case 10)
Review/approval gate for any agent output	**Ad hoc, tool-by-tool (AI notetakers etc.)**	Ongoing assessment of clearing and sanctioning tools, with reason and future integration into the knowledge platform

Process	Status today	What's needed
Culture/communication process	**Standard org comms and regular outputs**	As the AI strategy represents a culture change process, not just a technology rollout, the implementation plan reflects the change management options outlined
Peer review	**No structured approach exists**	Process with relevant owners of quality to establish quality standards for AI-based coaching and how these link into prioritizing human review

Implementation of technologies

The following steps are needed for the implementation of technologies:

- Platform setup with model/hosting confirmation and admin configuration.
- Knowledge base build after curation and metadata design, ingest curated subset; confirm no unconsented personal/participant data is included.
- Folder Sync configuration for the subset of material that changes regularly, connected only to approved, curated folders within Drive.
- First Tier 1 agents for KM retrieval, outputs tagging agent, horizon scanning, methods/writing coach agent. Each pilot follows the MVP-first principle (2–6 weeks) before scaling.
- Adoption-alert workflow (org-KPI 4) through a search-and-summarize pattern.
- Governed meeting-notetaking capability to replace ungoverned individual tool choice (Otter.ai, Read.ai) with either the platform's own capability or a consent-checked agent workflow.

Training needs

Training needs, coordinated with and led by the People team, follow a learning pyramid model (basic training for all, tiered up to AI values ambassadors, with more focused training for a small leadership group) rather than a single all-staff course:

- Baseline, all-staff: what The platform is, what it can and cannot see, the do/don't guide, and the sIfA AI disclosure process
- Prompt/task-specification training (COSTAR-style structured prompting) for anyone configuring agents or workflows
- Review AI literacy training: how to apply critical thought to all AI outputs through fact and sense checking and identifying bias
- Literature review methodologies
- AI values ambassadors: a small number of staff per team (research, ops, BD) with dedicated time to co-build and troubleshoot configurations to support peer learning
- Goal owners: Usage fluidity
- Training for the AI culture lead, knowledge worker, technology expertise specifically: curation/metadata judgment, ethics, governance, legal checks needed when curating the knowledge base.

Phasing and roles in implementation (35 months)

- Preparation (Months 1–8, February–September 2026): Capacity building/ structured training to develop an AI strategy, mandate clarification, evidence gathering, consultations, drafting of AI strategy, AI strategy certification, trialling platform, investment decisions, start of AI culture lead training
- Infrastructure building (Months 9–14, October 2026–March 2027): Project group recruitment (including knowledge worker and refining SoW for AI culture lead), six evidence-generating projects, drafting AI Literate Critical Thinking syllabus, first annual publication of public summary of AI work in Tafakari

- Pilot and continued infrastructure building (Months 15–20, April–September 2027): Knowledge base work continues with growing curated subset; first agents (use case 6 KM retrieval, one Tier 1 use case) piloted with one team; the six evidence-generating projects launch artefacts; change management processes continue. Infrastructure part 2: literature review processes, databases and methods coaching
- Scaling (Months 21–26, October 2027–March 2028): Expand agent/workflow coverage across remaining Tier 1 use cases; roll out training tiers organization-wide. Piloting literature review. Second annual publication of a public summary of AI work in Tafakari.
- Full integration (Months 27–33, April–October 2028): Agent/workflow coverage embedded as standard practice across all Tier 1 use cases; training tiers complete organization-wide; literature review moves from pilot to routine use. Review roll-out evidence year 1 (since infrastructure building) against org-KPIs 1/4/5.
- Next-phase decision (Months 34–35, November–December 2028): Decide the use case slate for the next phase and integration with next Busara org strategy, including whether Tier 2 use cases are warranted and whether the AI culture lead role should be continued; set the mandate and investment case for that phase. Third annual publication of a public summary of AI work in Tafakari.

Risks and quality assurance

The AI strategy faces ongoing risks in both development and implementation. Rather than listing them all here, they are listed in each section as they relate to governance, security and ethics, implementation, change management, and skills building.

- Project-delivery risks include milestone slippage
- Mitigation: through agile project management, a slipped milestone is used to assess whether capacity, capability or resource assumptions need adjustment before the whole implementation derails. Agility is maintained through project-management feedback loops on the obstacles and challenges each step encounters.

- Underestimating the effort it takes to curate the knowledge base.
- Mitigation: effort is already generously budgeted, if necessary, can draw on consultancy budget for additional support
- Quality is assured by testing each agent/workflow against defined review criteria before wider rollout; monitoring adoption and output-quality feedback via the named human reviewers assigned per use case; a structured feedback loop back into the AI culture lead, including through the change management events planned.

Our organizational context crystallizes our challenges. Busara is a decentralized organization, which poses a challenge when establishing central systems. To change this requires full commitment of the executive team to act as role models, AI strategy ambassadors and guardians of accountability. To support staff in committing to the AI vision, its values and guardrails requires building AI skill level assessments into performance reviews.

- A flattened AI voice is often appealing: we have worked hard to allow for a diversity of voices. However, AI is contributing to a narrowed set of expectations of what proposals or reports ought to sound like. To not fall victim to prioritising generic and polished prose, we will build out our writers development program and publication support.
- Skills take time: building skills takes time, as does maintaining them. To balance the time commitment of learning both AI skills and job-relevant manual skills while continuing our work, we offer a tiered training program and prioritize a few processes that are likely to pay off in saved time swiftly, such as use cases 5 (project management) and 6 (knowledge management)
- Trade-offs can bear a reputational risk: being transparent can get us dismissed. Saving time by using AI can damage our reputation. Increasing our environmental footprint to do our work quicker is hypocritical while also working on climate change. This requires us to weigh the costs of trade-offs, reputation and principle through transparency and by drawing clear lines of what we do not do because the trade-off is too high.

References

1. Behavioural Research UK. 2026. *Using Generative AI for Behavioural Research Effectively and Responsibly: a BR-UK Statement*. London: BR-UK.
2. Bender, Emily M., and Nanna Inie. 2026 (7 January). "We Need to Talk About How We Talk About 'AI'." *Tech Policy Press*. https://www.techpolicy.press/we-need-to-talk-about-how-we-talk-about-ai/
3. Bowman, AI 2025. AI in the Loop.
4. Busara. 2026. *Busara Strategy 2026-2028: Practice Partnerships for Impact*. Nairobi: Busara.
5. Center for Global Development, The Agency Fund, and IDInsight. 2026. *Generative AI Playbook for the Development Sector*. Washington, D.C.: Center for Global Development.
6. Centre for International Governance Innovation. 2026. *Human Analysis Standard: Guidance for Responsible AI Use in Policy Research*. Waterloo: Centre for International Governance Innovation.
7. European Parliament and Council. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) [2024] OJ L 2024/1689.
8. Hao, Karen. 2026. *Empire of AI: Inside the reckless race for total domination*. New York, NY: Penguin.
9. ISO. 2023. *International Standard ISO/IEC 42001 (First Edition): Information technology — Artificial Intelligence — Management system*. Geneva: ISO.
10. Ministry of Information, Communications and the Digital Economy. 2025. *Kenya Artificial Intelligence Strategy 2025-2030*. Nairobi: Ministry of Information,, Communications and the Digital Economy.
11. Muange (Manager Project Operations and Busara ISERC), Gladys. June 2026. *Busara Risk Management Policy*. Nairobi: Busara.
12. Muange (Manager Project Operations and Busara ISERC), Gladys. May 2026. *Busara Safeguarding Policy*. Nairobi: Busara.

13. Muange (Manager Project Operations and Busara ISERC), Gladys. September 2025. *Data Protection Policy From Data Collection to Data Storage*. Nairobi: Busara.
14. Muange, Gladys. March 2026. *Research Ethics Policy : Ethics and Compliance at Busara*. Nairobi: Busara.
15. Omedo (People Officer IT Systems and Support), Patricia. November 2024. IT *Operations & Systems Policy*. Nairobi: Busara.
16. Saleh, Engy, and Mareike Schomerus. 2026a. *Busara Futures: Org-Wide Survey and Consultations Report (internal document)*. Nairobi: Busara.
17. Saleh, Engy, and Mareike Schomerus. 2026b. *The sIfA tool for a Statement of Intellectual Fellowship and Accountability: An invitation to reflect on how humans and AI interact in producing knowledge (Groundwork 32 Thought Piece)*. Nairobi: Busara.
18. Sampson, Rob , and Anmol Singh Rawat. 2025. *Lessons from IDinsight's Generative AI transformation:* ID Insight.
19. Schomerus, Mareike. 2026. *sIfA Tool — Statement of Intellectual Fellowship and Accountability (Version 1.4) [Computer software]*. Nairobi: Busara. www.sifatool.org.
20. The Republic of Kenya. 2019. *Data Protection Act CAP. 411C (Published in Kenya Gazette Vol. CXXI—No. 156 on 15 November 2019, Assented to on 8 November 2019, Commenced on 25 November 2019, [Revised by 24th Annual Supplement (Legal Notice 221 of 2023) on 31 December 2022]*. Nairobi: The Republic of Kenya.
21. The Republic of Kenya. 2026. *The Government of Kenya: The Artificial Intelligence Bill. 15/4, Kenya Gazette Supplement Senate Bills.*
22. West, Robert, and Susan Michie. 2021. "A Brief Introduction to the COM-B Model of Behaviour and the PRIME Theory of Motivation." *Qeios*. doi: 10.32388/WW04E6.3.
23. Yalouris, Anna. 2026. *Strategy, Vision & Mandate for Partnership Development Vertical (draft for comments)*. Nairobi: Busara.

Appendices

Appendix 1: PROTECT Code of Busara's SCP

SCP's PROTECT code could serve as a starting point for an initial do this/don't do that guide on using AI.

P — ***Proportionate use:*** use AI freely for low-stakes, reversible, internal tasks (brainstorming, first-pass literature screening, meeting notes, code debugging). Apply much greater caution for high-stakes outputs, anything shared externally, anything used to inform decisions, or anything representing Busara's expertise to clients or partners. The higher the consequence of an error, the more human judgment should dominate.

R — ***Review it, always:*** only use AI to support tasks you already know well enough to spot errors in. If you can't judge what "good" looks like, you can't catch what AI gets wrong.

O — ***Own the decision:*** the rationale for any methodological, analytical, or strategic decision must come from the staff member's own reasoning, never "AI said so." AI generates options and surfaces considerations; the judgment call and its documented rationale stay human.

T — ***Transparency:*** always disclose AI use, including for internal tasks like screening or brainstorming — a simple note is enough. This directly targets the current culture of silent rewriting and undisclosed distrust identified above.

E — ***Expertise:*** protect the specific skills that make each role valuable (qualitative analysis, critical interpretation, argument-writing, fieldwork problem-solving) by deliberately not outsourcing them to AI, even when AI would be faster — productive struggle is how expertise is built and maintained.

C — ***Confidentiality:*** never input participant data, client information, proprietary research instruments, or partner data into a public AI tool. When in doubt, anonymize or don't use AI. Always check client/partner policy before using AI on meeting recordings.

T — ***Trust but verify:*** every AI-generated factual claim, citation, or statistic must be independently verified before use; AI fabricates citations and produces confident, plausible-sounding errors. Running multiple prompts and treating disagreement between outputs as a signal to investigate is a useful check.

Other guardrails that have been suggested are:

- In research, AI is used only for checking and challenging (data analysis review, framing critique, research design feedback) via direct researcher prompting — never for initiating analysis — with this guardrail stated in every project's methodology.
- In writing, AI is not used to produce first drafts of core documents; it supports only reviewing for critical feedback and editing, protecting writing as thinking.

Appendix 2: Mandate

This AI strategy was formally commissioned by Busara's CEO Francis Meyo with full authority for its development delegated to the author Mareike Schomerus, VP Futures, Impact, Voice and Engagement, who was supported by Engy Saleh, Director of Behavioral Research and Academic Engagements (BRACE). It was developed through a process of external training of the author, internal consultations, research and peer feedback. It is a living document, but will be signed off for completeness and feasibility by the CEO, after a final organizational feedback round on the full strategy.

Appendix 3: Methods used in developing this AI strategy

To deliver the AI strategy first required capacity building of the project lead Mareike Schomerus (VP FIVE), staff time of the co-lead of the Futures Initiative Engy Saleh (Director BRACE), consultations space (org-wide meetings, consultation sessions, surveys), and executive team commitment and input. The timeline for strategy development is in Table 11 on the following page.

Table 11. Activities plan for AI strategy development

Q1	Q2	Q3	Q4
Start of certified AI Manager training for strategy author, consultations	Training extended, first public good contributions on AI transparency (sIfA), consultations	AI culture lead training, hands-on platform try-outs, delivery of AI strategy; certification	Implementation as per implementation plan

1.1.2 Part 1: Setting the scene: consultations about where the organization is

- Workshop with knowledge management team (IT, software, compliance) (February 26)
- Conversation with the research excellence team on whether knowledge management is a good pain point to prioritize (Feb 5, 2026)
- Org-wide consultations (average show-up of 15 staff members, one outlier meeting with only 2 people) (March 10 in person, March 12 in person, March 17 online, March 26 online)
- Org-wide survey: Shape how Busara engages with AI (n=51)
- Output: State of AI in Busara report (see 16.8 Appendix 8: State of AI in Busara report (separately appended).
- Slack search on informal governance recommendations on AI

1.1.3 Part 2: Socialising first steps

- Executive team: mandate development session 1 (April 16) — current status of AI in the world and in Busara, technology, governance, broader strategic directions for Kenya, typical pitfalls and great opportunities, AI usage in Busara and its implications (findings from survey and consultations), pitfalls of AI usage, developing the strategy, mandate development and expectation management, quick wins: sIfA
- Executive team: mandate development session 2 (May 12) — developing Busara's AI vision

- Executive team: mandate development session 3 (May 28) — deep pain and use cases, values, AI strategy mandate, AI strategy canvas, critical needs, decisions needed, next steps
- Shaping Busara's AI contribution approach (online, 30 April) — to get feedback on the approach and design of the sIfA tool (with FIVE, Research Excellence, RAP, design, software)
- Additional sense-checking of sIfA with design and software departments

1.1.4 Part 3: Deepening the understanding

- Interviews and evidence gathering with operations as well as IT, software, ethics and compliance, people and finance teams
- Org-wide AI consultation Session 1: what is an AI strategy? Developing Busara's vision for the future (June 2 and June 3)
- Org-wide AI consultation Session 2: Understanding processes and values in Busara's AI strategy (2 July, 3 July)
- Total n across both consultations: around 40

1.1.5 Part 4: Try and learn consultation sessions

Hands-on 3.5 hour experimentation with use cases 1 and 6 on the platform; 4 repeated sessions (August 12/13) for three different time zones (India/Latin America/East Africa) with the following agenda (11 participants — 7% of staff).

Agenda for the hands-on session — goals of the session:

- To familiarize with some of the thinking behind the AI strategy
- To familiarize, learn and give feedback on the platform
- What does it mean to work with a model-agnostic platform (what do different models produce/what models might we want to reject)?
- To work with value-based Socratic pushback vs 'ask a question, get an answer' approaches to build critical thinking skills
- To trial metadata curation, tagging and knowledge base retrieval (using mock knowledge base with real Busara examples)
- To see what we are learning: What is encouraging? What is concerning?

Author training and feedback opportunities

The author of this strategy attended a structured training towards becoming a certified AI manager with TÜV Austria (32 teaching units). This was further supplemented with additional seminars and webinars on AI in research, for example with IDInsight, BR-UK or the LSE, and a substantive amount of autodidactic upskilling. The author is now attending additional training to become an AI culture lead with the Happy Company. This training, spread over 18 months, is designed to accompany AI leads in implementing AI strategies, with a particular emphasis on ethics, governance, culture and technologically-driven change management.

After developing the sIfA tool, the author further sought feedback on the tool from relevant communities by launching it publicly at the On Think Tanks conference in Rabat (May 2026), the BR-UK launch of the statement of responsible use of generative AI in behavioural research (May 2026), with further talks scheduled.

Survey instruments for org consultations

Org-wide survey on Busara Futures: Shape how Busara engages with AI

Busara is taking a thoughtful look at how AI fits into our work—today and in the future. Before any decisions are made, we want to hear directly from you. Not just what tools you use, but how you actually experience AI: what excites you, what frustrates you, what feels unclear, what feels wrong, and what you wish worked better.

The following questionnaire is completely anonymous, and there are no right or wrong answers. Whether you use AI daily, avoid it entirely, or are still figuring out what it even means for your role, your perspective is valuable. Please feel free to be candid and as detailed as you'd like; honest responses (including doubts and concerns) are far more helpful than polished or "expected" ones.

Your input will directly inform how Busara approaches AI—what we encourage, what we're cautious about, and how we can support staff in using AI responsibly and effectively. This is a chance to help shape a shared approach that reflects our realities, and the way we actually work. Thank you for taking a few minutes to contribute to this conversation.

- Please describe what AI is to you (in general, in life, in work...etc)
- What are you most excited about when it comes to AI? (generally)

- What are you most fearful of or anxious about when it comes to AI? (generally)
- Familiarity with AI: division affiliation; familiarity ratings (scale of 4) for ChatGPT, Claude, Gemini, CoPilot, Claude Code, OpenClaw; other AI tools used and familiarity level.
- AI Use at Work: whether AI is currently used (scale of 4); how it is used (drafting/editing, re-write option, proof read option, summarizing, meeting notes, brainstorming, research support, data analysis/coding/debugging, presentations/visuals, social media posts, learning/upskilling, checking own work, project management, other); which tools and platforms are used.
- Opportunities and Pain Points: what parts of work feel most mundane/repetitive; which challenges AI could most help reduce (administrative/documentation burden, proposal writing, knowledge management/evidence synthesis, project management, time management, data cleaning, data analysis, writing/editing, presentations/visuals, literature reviews, other); which cognitive skills to protect (data analysis, problem formulation, seeing through a lot of information, making connections, coming up with new ideas, writing in own voice, other); how AI could support knowledge management/evidence synthesis and proposal development; what perfect AI would free people up to focus on.
- Good Use of AI: confidence in using AI ethically and responsibly (1-5); what constitutes responsible and ethical use; concerns (data privacy/confidentiality, accuracy/reliability, bias/ethical risks, over-reliance, unclear guidelines, deskilling, streamlining of knowledge, own work/prompts/data feeding the model, no concerns, other), with room to elaborate; what would help feel more comfortable using AI (clear guidelines/policies, training/examples relevant to role, approved tools/use cases, opportunities to experiment safely, ongoing discussions/support, other).
- Final Thoughts: anything else to consider; how AI-related questions should be debated at Busara (webinars, dedicated meetings/office hours, involvement in trial runs, other).

Org-wide survey: Identifying AI use cases (also adapted for direct input from all executive team members)

- Who are you? Are you happy for us to follow up?
- Your goal and processes: choose a goal (one objective, linked to org strategy); identify the processes required to achieve it.

- Skills and learning: most important skills for this goal; moments in the process where the most learning happens.
- Identifying the deep pain: which challenge this goal faces; which parts of the process are challenging; will the challenge grow bigger and what would make it grow; are mitigating approaches already in place; what happens if no action is taken.
- What could a tool look like? Imagine a perfect solution: who uses the tool; what problem it solves; what it produces; where the data comes from.
- Trusting your tool — what would make you trust it: outputs grounded in Busara's organizational knowledge; transparent on all sources; AI use clearly declared; honest about its limits; teaches rather than just does the work; gives feedback on where to improve; outputs safe and confidential; writes in Busara's voice; outputs aligned with Busara's positions; stops when asked something not okay to do; audit trail for all actions; a named accountable person; clear human role (human-led/in-the-loop/on-the-loop/automated); other.

Org-wide survey: setting context for AI

- What is my primary work and responsibility? How skilled am I at this (beginner/intermediate/advanced, explained in own terms)?
- How do I work? What tools do I use, for what, and where are these? How do I save my work?
- What is my ideal way of interacting with AI? (e.g. concise instructions, strong guard rails, more/less explanation, level of autonomy)
- Where am I seeking AI help? (list use cases)
- What style do I prefer? How should AI talk to me? What guardrails do I want in place?

Org-wide survey: Seeking knowledge at Busara: what does it take?

This form asks about how much time you spent on searching for information. What level are you? What department are you (BizOps/RAP/MRI/FIVE)?

- Think of the last time you needed a past instrument, finding, publication or proposal you knew existed somewhere. Best case, how long did it roughly take to find it (10 minutes/30 minutes/60 minutes/more than 60 minutes/other)? Worst

case, how long has a search like that taken (more than 60 minutes/more than half a day/usually give up after first attempt/other)?

- Redoing work: in the past six months, did you redo work (rebuilding an instrument, re-running analysis, rewriting a standard paragraph) because you could not find the original? If yes, how long did redoing it take, and roughly how often did this happen? Can you describe a recent example of what you were looking for, where you looked, and what you found or didn't find?

Appendix 4: sIfA

The sIfA tool (www.sifatool.org) records high-level and detailed information on human/AI interaction in the production of a piece of work, using the Contributor Role Taxonomy (CRediT) by default. Audit trail (all notes provided by Claude Sonnet 5/Opus 5, logged after each session).

Investigation

27 Jul 2026: analyzed the newly-provided Busara 2026-2028 strategy document ('Practice Partnerships for Impact') against the TÜV ten-segment structure and the expanded five-KPI framework; findings (ten-segment mapping, KPI-to-AI-segment mapping, full source text) appended to strategy-gathered-evidence.md, closing the item-7 'no standalone strategic plan found' gap.

4 Aug 2026: grouped the five org KPIs under synthesised strategic-pillar headings and mapped her six Segment-1 goals against them; findings (goal-to-KPI table, KPI grouping rationale, the unbuilt-toward KPI 3/KPI 4 gap, and a timeline-asymmetry note on KPI 5 vs. KPI 4) appended to strategy-gathered-evidence.md as a new 'Goal-to-KPI mapping' section.

28 Jul 2026: checklist cross-check on the same draft surfaced new partial evidence for checklist items 10 (client/funder AI clauses — current contract and RFP landscape) and 18 (participant-data consent scope — a policy conclusion that all data collected to date is currently unusable for internal AI-assisted querying); both logged into strategy-gathered-evidence.md without marking either item resolved.

4 Aug 2026: analyzed the Slack search she ran for item 8 (informal AI guidance) and organised findings by segment — updated item 8 (an ai-adhoc-working-group channel existed with real unfinished work, including a planned but unconfirmed 'Values Statement on AI'), added a change management example (Segment 3), a knowledge-management AI risk framework and a trial-platform data-scope fact (Segment 5/6), Kenya DPA consent/anonymisation reasoning adding detail to item 26 (Segment 8), her own Gen AI Evaluation Playbook critique as near-ready Segment 2/9 material, and a new KPI 5 baseline figure (256 outputs/3.5 years). Updated strategy-information-gathering-checklist.md item 8 and strategy-master-todo.md's Phase 1 checkbox to reflect the Slack search is done.

16 August 2026: ran a full quality check against the TÜV ten-segment rubric and the 14 August outstanding-issues checklist; corrected several internal inconsistencies introduced by an earlier restructuring; assembled a paragraph for the Executive Summary at her request, using only verbatim sentences already existing elsewhere in the document; re-verified the document after she made further edits independently

Validation

27 Jul 2026: two-pass rubric feedback (TÜV guidance + Woche-8 exemplar standards + gathered evidence) and checklist cross-check on her Layer 0 mandate draft — flagged the empty SMART table, a function-name inconsistency ('FIVE' vs. the strategy document's 'Voice, Impact & Futures'), and thin Prerequisites/Activities tables; logged a new partial signal on checklist item 3 without marking it resolved.

28 Jul 2026: two-pass rubric feedback (TÜV guidance + Woche-8 exemplar standards + gathered evidence) on her Layer 0/Layer 1 draft (AI strategy.docx) — flagged production issues (an accidentally-pasted raw AI-interview transcript, two verbatim checklist-file copy-paste fragments, a malformed citation, an orphaned placeholder heading), confirmed the SMART table and Activities-plan/pilot-group tables remain thin since 27 Jul while the FIVE/Voice-Impact-Futures naming inconsistency is now resolved, and verified all eight standing 'strategic tensions' from strategy-gathered-evidence.md appear somewhere in the draft.

4 Aug 2026: compared the five staff-sourced values in section 2.8.1 of her rough draft ('AI strategy rough draft section 0 and 1.docx') against the four AI-strategy values and six goals in the Section 1 draft. Found four of five values closely matched, two nearly word-for-word (including a shared typo confirming direct textual reuse), and flagged one value with no match anywhere in the AI strategy ('Openness to challenge'), plus a related gap in 'Thinking' ('dealing with contrasting views'). Also flagged a tone difference: the staff-values section admits the values are only sometimes lived up to; the AI-strategy values are stated purely aspirationally. Findings appended to strategy-gathered-evidence.md as a new 'Values comparison' section.

4 Aug 2026: built a full goal-to-Org-KPI table listing every measurable in all six Section 1 goals against the five org KPIs, at her request. In building it, found Goal 3 and Goal 4 both now contain a version of the KPI 4 adoption-alert mechanism, inconsistently: Goal 4's version keeps the human-validation safeguard and is correctly scoped to the pilot cohort; Goal 3's version conflates internal template adoption with partner adoption of Busara's methods, drops the safeguard, and contains a garbled phrase ('to support Google citations'). Recommended keeping the mechanism in Goal 4 only. Also noted Goal 1 was updated to explicitly name 'openness to challenge,' closing the gap found in the values comparison. Findings appended to strategy-gathered-evidence.md as a new 'Goal-to-Org-KPI table' section.

6 Aug 2026: reviewed her Segment 3 (Stakeholders & change management) draft against the three-layer rubric (TÜV guidance, Woche 7/8 exemplar standards, gathered evidence) plus a checklist cross-check; flagged the blank Einfluss×Interesse stakeholder table, a left-over generic template table, department-title mismatches against her own sourced org chart (MRI vs. Research Excellence; BRACE lead title), the SLT-only 'team ambassadors' vs. her own dissenting-voices commitment, and an unresolved ISERC/item-26 ethics-infrastructure question; logged partial advances on checklist items 3 and 32 into strategy-gathered-evidence.md; compared against the sealed Claude draft's Segment 3 (unsealed 24 July) and reported the differences directly to her.

6 Aug 2026: reviewed her Segment 7 (Solutions & resources) draft against the three-layer rubric (TÜV guidance, Woche 5/8 exemplar use-case evaluation framework, gathered

evidence) plus a checklist cross-check; flagged the segment's word count alone (~2,550 words) equalling the whole document's ~2,500-word target, the missing explicit use-case suitability/risk/complexity/cost scoring and HitL/HotL/autonomous labeling the exemplar rubric calls for, a second independent occurrence of the MRI/Research-Excellence department-name mismatch first flagged in the Segment 3 review, two of five consultation-sourced AI-free zones missing from the scope exclusions, and the persistent environmental/compute-cost silence; logged new vendor specifics into strategy-gathered-evidence.md item 13 (advancing checklist item 13) and confirmed the segment correctly closes the data-privacy silence and honors the ≤2-expert-hires non-goal.

6 Aug 2026: at her request, cross-checked Segment 7's Tier 1/Tier 2 use cases against the six Segment 1 goals (via the existing Goal-to-Org-KPI table) and Segment 2's value proposition. Goals 3, 4 and 6 strongly served, Goals 2 and 5 thinly; identified that Goal 1 (AI-literate critical thinkers) has no supporting use case for its one concrete KPI-5-linked deliverable, the quarterly AI/social-science synthesis publication, which also maps to a named Futures capability ("AI Research Agenda") in the sourced corporate strategy. Flagged two secondary consistency notes (Goal 4's vector-base target sitting in Tier 2 rather than Tier 1; Goal 5's root-cause-classification language unmatched to any use case output). Confirmed no gap between Segment 2's two anchor KPIs and Segment 7's coverage of them. Findings appended as an addendum to the 6 August Segment 7 review in strategy-gathered-evidence.md.

8 Aug 2026: two-pass rubric feedback (TÜV guidance + Woche-8 exemplar standards + gathered evidence) plus sealed-draft comparison on Segments 6, 8 and 10, submitted across one session. Segment 6 (Data strategy): flagged a likely KM-backlog page-count error against the file's own doubly-independent; praised the granular 13-source data inventory and the use-case data-requirements crosswalk. Segment 8 (Security & ethics): logged two new facts meaningfully advancing checklist item 26 (NACOSTI IRB-granting accreditation; Director of Operations owning data protection), confirmed item 29 correctly applied; flagged the missing Antony consent-precedent example, a possible jurisdiction-coverage overclaim, and unsourced PROTECT/BenevAI attributions. Segment 10 (Revenue & cost efficiency): resolved checklist items 24 and 25 outright via real finance/ops interview data now sourced in the draft, resolved item 34, verified all arithmetic across

the platform costing/repetitive-task/KM-ROI/staffing tables (all reconciled), and caught two substantive errors — a KPI 2/KPI 3 label swap on the margin scenarios, and a ~15x understated break-even claim in the cost-efficiency paragraph. Also flagged, now a confirmed cross-segment pattern: "MRI" recurs a third and fourth time (Segments 8 and 10) where the sourced org chart names the department "Research Excellence" — recommended a single global fix.

6 Aug 2026: reviewed her Segment 5 (Technologies & tools) draft against the three-layer rubric (TÜV guidance, the Woche 1/2/5 + Hardware exemplar framework — infrastructure ladder, L1-L4 solution levels, agentic surcharge, MVP-first phasing, FinOps for token costs — and gathered evidence) plus a checklist cross-check. Credited strong evidence-consistent reasoning (infrastructure-ladder logic, LangGraph-vs-MemGPT agent-autonomy argument, correct direct use of the previously-flagged Otter.ai risk detail) even where the exemplar's own vocabulary wasn't used; flagged an unsourced IDinsight platform-consolidation claim, a forward-looking tension between her EU-hosting choice and Kenya's own data-localisation/residency regulatory direction (Bill cl.36(2)(e), national AI Strategy Pillar 2), and a concrete gap on paper-format backlog material having no tooling path. Compared against the sealed Claude draft's Segment 5 and reported the differences (her single-platform choice more disciplined than my assumed multi-tool stack; my placeholder "already licensed" claim was wrong; neither draft addresses FinOps/token costs).

6 Aug 2026: reviewed her Segment 9 (Industry-specific applications) draft against the three-layer rubric and a checklist cross-check. Confirmed as factually accurate: the sIfA/CRediT claims (14 standard roles, 0-2 extent scale, checked directly against sifa.schema.json), the Global Majority differentiation hook, the Kenya AI Strategy priority-sector convergence, and the Kenya AI Bill high-risk-sector list. Flagged an unsourced IDinsight internal-transformation statistic (a second, separate unsourced IDinsight claim alongside Segment 5's), a cluster of research-integrity/governance bodies new to the evidence base, the missing Global South Network for Trustworthy AI / Masakhane citations the market scan ties directly to this exact claim, and two of her own five Kenya-strategy points (localisation, local-talent-development) being absent. Compared against the sealed Claude draft's Segment 9 and identified a factual error in my own sealed draft: it claimed IRB integration into AI governance as an existing Busara differentiator, which item 26 directly contradicts (no ethics/IRB infrastructure exists) — corrected the record rather than leaving it unflagged.

8 Aug 2026: two-pass rubric feedback (TÜV guidance + Woche-8/Woche-4 exemplar standards + gathered evidence) plus sealed-draft comparison on Segment 4 (Implementation plan), plus a requested cross-check of Segment 4 against Segment 1's Goal-to-KPI table and Segment 7's use-case catalog to confirm every KPI/goal deliverable and use case links into implementation. Findings logged as-is at her request, with an explicit caveat that the cross-check runs against the last-reviewed Segment 7 (6 August, eight Tier-1 use cases) rather than Segment 6's newer ten-use-case list, so use-case-numbering findings are provisional. Flagged two numbering mismatches (Segment 4's "use case 10" and "use case 7" references don't match Segment 6's definitions), two deliverables with no visible implementation step (the quarterly AI/social-science synthesis publication tied to Goal 1/KPI 5 and use case 10 — a continuation of a gap first flagged in the 6 August Segment 7 review; Goal 2's annual Tafakari public summary; use case 8/literature screening), and two thinner rubric elements (unlinked project-group ownership, risks without paired countermeasures). Noted a genuine strength: her own text already self-flags the AI-budget-forecasting gap for Goal 3 rather than glossing over it.

8 Aug 2026 (later session): full-draft review of the complete Aug 8 strategy document (all ten segments) against the official TÜV four-point rubric and against the AI Transformation Checklist (all eight sections, A–H — she pasted A–C, the full A–H source file was located and used). Found and reported: the Executive Summary and Bibliography are still placeholders; the six-goal 'Quantify the Value Proposition' section is an unfilled template repeated verbatim per goal; the Interests/Influence stakeholder-map table (29 rows) is empty; 'MRI' is defined two different ways in the same document; Goal 2's own Tafakari-publication timeline (Q4 2027) contradicts the Implementation Plan's phasing table (Months 8–10); Use Case 3 (proposal/BD support) is claimed by both Goal 3 and Goal 6 without reconciliation; compiled a full abbreviations/acronyms list. At her explicit direction (a structural/mechanical request, not content drafting), also executed two edits directly in the docx: (1) fixed a white-on-black table by changing 45 cell-shading values from near-black to white, no text touched; (2) relocated the ten Tier-1 use cases from the Solutions & Resources section to sit under their assigned goal (mapping decided together, including three genuinely ambiguous cases she resolved directly), renumbered them 1–10 in the new goal order, and propagated the renumbering to the parallel Data Strategy use-case listing and ~20 scattered cross-references elsewhere in the document. Backed up before each edit; validated paragraph counts and OOXML schema before and after.

8 Aug 2026 (same session, follow-up): filled the Executive Summary placeholder ('tk') at her explicit direction, but strictly as a cut-and-paste assembly, not drafting — zero original prose. Pulled five verbatim excerpts already existing elsewhere in the document: the AI Vision Statement sentence (labelled 'Ambition'), the five bold-step titles, the six goal headings, the ten (newly renumbered) use-case titles, and the five named challenge labels from the 'what hinders us' section. Confirmed with her first that this counted as assembly rather than an exception to the review-only/no-drafting arrangement on this project.

14 Aug 2026: full-document quality check on KPI, measurement, implementation-plan and alignment.

14 Aug 2026: verified cross-reference integrity by resolving REF field bookmarks rather than reading displayed text. Found a mis-targeted field in the security appendix — '(see Appendix 4: PROTECT Code…)' pointed at Appendix 3's bookmark while displaying Appendix 4's title, so a routine field update would have silently rewritten it to the wrong appendix — and repointed it. Confirmed by the same method that the 'section 3.14' reference needed no edit, its bookmark being correct and the number regenerating on update. Every edit XSD-validated against the pre-edit backup with paragraph counts reconciled.

14 Aug 2026: reconciled Table 4 with the Goals section and the vision statement across eight conflicting KPIs. Four resolutions were hers (conference tool 75%; a single AI-assessed literacy metric at 75%; sIfA scope 'substantive'; dropping the 30% RAP figure, which had no baseline anywhere in the document). Four were determined from the document's own logic (100% migration as end-state; 'all staff' denominator, since Goal 1's Specific line already said so; 'AI-searchable database' not 'vector base', since vector RAG is designated Tier 2; and reattributing an orphan proposals metric from Goal 4 to Goal 6). Also fixed a live typo in Goal 1's Measurable line ('75% how an AI-assessed increase' → 'show').

14 Aug 2026: corrected Goal 5's impossible 'Q4 2025' milestone to Q4 2026, then on her instruction moved the dependent 'Implement fixes' step from H2 2026 to H1 2027, which had overlapped and partly preceded the audit it depends on. Corrected the evidence-generating project count from four to six in two phasing entries — verified first against Table

15, which lists six projects with durations rather than start dates, so no deliberate subset of four was intended — and normalised 'evidence-gathering' to 'evidence-generating'.

14 Aug 2026: rebuilt the appendix structure — renumbered sequentially 1–15, eliminating a duplicate 'Appendix 3', two unnumbered appendices, and gaps at 7 and 8; numbered and retitled an orphaned Heading 2 as Appendix 13 and restored its displaced block label as an in-text heading. Repointed seven main-body cross-references that all read 'Appendix 9', and removed a false promise that the Key Values elaboration sat in an appendix (it is still in the main body at 3.4). Added a missing caption to the budget-overrun scenario table, preserving the SEQ field so all table numbers regenerate by position.

September 11 2026: Quality check for duplications, inconsistencies, formatting errors

Writing (reviewing and editing)

4 Aug 2026: two-pass rubric feedback (TÜV guidance + Woche-8 exemplar standards + gathered evidence) on Segment 1 (Vision and Goals) draft, plus the Phase-4 blind comparison against the unsealed Claude draft (seal lifted 24 Jul). Flagged word count (~4,150 words vs. the ~2,500-word whole-document target), missing explicit linkage of the six goals to leadership's five-KPI framework, missing nesting to the named 'Practice Partnerships for Impact' corporate strategy, and a scope question on whether Canvas-worksheet material is embedded in the narrative section. Checklist cross-check: no open items (3, 10, 16, 18, 19, 24, 25, 26, 27, 28, 34) resolved by this submission.

5 Aug 2026: two-pass rubric feedback (TÜV guidance + Woche-8 exemplar standards + gathered evidence) on her Segment 2 (Value proposition of AI) draft, plus the Phase-4 blind comparison against the unsealed Claude draft's own §2. First pass (six-goal draft): flagged that it restated Section 1's six goals rather than adding a value lens above them, that the risk/mitigation table belonged to Segment 8 not Segment 2, that neither KPI 1 nor KPI 5 (the two KPIs her own KPI-to-segment mapping ties to this segment) were named anywhere, an unsourced point-estimate figure (20% admin-cost reduction, checklist item 34 territory), a mislabeled 'Level of text processed in Goal 1' tag under what was structurally Goal 5, an empty prioritisation-matrix placeholder, and the opening paragraph's un-

addressed echo of the efficiency/impact framing tension against the corporate strategy's own 'growth as a consequence of impact' language. Second pass, after she revised to a single framing paragraph plus three perspective bullets (For Busara/For partners/societal benefit): confirmed the revision resolved the restatement, missing-KPI, and framing-tension issues; answered her question on how to name the KPIs by confirming the corporate strategy's own Section 10 does not number them (the 'KPI 1'/'KPI 5' numbering is her own analytical convention from the evidence-gathering work), recommending she source that numbering explicitly at first use rather than let it read as an official Busara label, and flagging that 'learning assets (KPI 5)' should be tightened to the KPI's literal name (Evidence and Learning Asset Rate). Checklist cross-check: no open items (3, 10, 16, 18, 19, 24, 25, 26, 27, 28, 34) resolved by either version.

8 Aug 2026 (same session, follow-up): fixed the Tafakari annual-public-summary date conflict flagged in the full-draft review — Goal 2's Time-bound line said Q4 2027, the Implementation Plan phasing table said Oct–Dec 2026 (Months 8–10). She confirmed Oct–Dec 2026 is correct; changed Goal 2's line from "Q4 2027: Publish first annual public summary." to "Q4 2026: Publish first annual public summary." Single, unique string match; backed up before editing; validated after.

15 August 2026: condensing pass; typos and broken syntax corrected; reconciled Table 4 with the Goals section and the vision statement across eight conflicting KPIs; corrected the evidence-generating project count from four to six in two phasing entries; wrote connective prose in executive summary; full-document quality check on KPI, measurement, implementation-plan and calculation alignment, draft Kanban board.



Busara Groundwork
THOUGHT PIECE

About Busara

Busara is a research and advisory organization, working with researchers and organizations to advance and apply behavioral science in pursuit of poverty alleviation. Busara pursues a future where global human development activities respond to people's lived experience; value knowledge generated in the context it is applied; and promote culturally appropriate and inclusive practices. To accomplish this, we practice and promote behavioral science in ways that center and value the perspectives of respondents; expand the practice of research where it is applied; and build networks, processes, and tools that increase the competence of practitioners and researchers.

About Busara Groundwork

Busara Groundwork lays the groundwork for future research and program design. As think pieces, they examine the current state of knowledge and what is needed to advance it, frame important issues with a behavioral perspective, or put forward background information on a specific context.

How to cite:

Schomerus, Mareike. *Busara's AI strategy 2026 – 28: Skilled, value-driven, and equitable AI interaction for stronger research, advisory, and practice. Busara Groundwork 36.* Nairobi: Busara, 2026. DOI: doi.org/10.62372/HIDB1128

38 Apple Cross Road,
Lavington, Nairobi, Kenya
www.busara.global

Busara