

Wakamoso Africa

Ethics Framework

The institutional rubric every tenant questionnaire is reviewed against — refreshed May 2026

Version 1.2.1

Effective date: pending adoption (refreshed 8 June 2026).

Owner: Information Officer

Information Officer: Mélanie Prinsloo (CEO). Deputy Information Officer: Michael Matthews (CIO).

IR Registration Number: 2026-018959 (registered 25 May 2026).

Internal counsel review: Malusi Mazibuko.

Filing: Governance / POPIA Compliance

1. Purpose

This framework is the institutional standard against which every questionnaire fielded on the Wakamoso Platform is assessed before it reaches any data subject. It implements the Ethics Review Office obligation in the Subscriber Agreement v3.0, is intended to support the 'appropriate safeguards' requirement of POPIA Section 32 (research exemption framework), subject to regulator interpretation, and operationalises Wakamoso's ethics-by-design positioning.

It applies to questionnaires fielded on the Wakamoso Platform, regardless of topic, tenant type, or research method, in the ordinary course. A municipal cost-of-living tracker, a corporate workforce engagement survey, an academic baseline, a customer-satisfaction study, a loyalty-programme research instrument, and a C3B community capability baseline all flow through the same framework. The depth of review scales with the risk profile of the questionnaire content.

v1.2.1 reflects calibration learnings from EAR-2026-0001.01 through the location-regime decision (DEC-2026-002). The operative rubric is maintained in the Governance Dashboard (lib/ethics-assessor/framework.ts, FRAMEWORK_VERSION v1.2.1); this document is the governance description of that rubric and is refreshed to match it.

Calibration history (v1.1 to v1.2.1):

v1.1.1 - Section 26 closed-list calibration (aggregation of non-special data does not create special-category data); C14 (cumulative profiling) and C15 (retention and lifecycle) added as Core checks.

v1.1.2 to v1.1.4 - C5 Opt-In alignment calibrations (inbound-as-opt-in; recognised-category auto-pass) and C2 scope restoration.

v1.1.7 - consent-gate false-premise fix: there is no per-questionnaire consent gate; consent is the platform POPIA Opt-In v2.0.

v1.1.10 - severity discipline (proportionate findings).

v1.1.13 to v1.1.14 - C4 surface architecture: notification at collection is Surface 3 (the questionnaire body intro), distinct from the canonical Opt-In and the entry path.

v1.1.15 - POPIA / ethics domain bifurcation: POPIA-domain findings bind the disposition and may emit conditions; ethics-domain findings (C10, C11) are advisory.

v1.1.16 to v1.1.17 - anti-invention and surface-vocabulary discipline.

v1.2.0 - mobile-app delivery modes added.

v1.2.1 - C9 location-regime calibration (DEC-2026-002).

2. Structure

The Framework has three layers:

- Layer 1 — Scenario triage (Section 3). At the start of each assessment, the Office classifies the questionnaire as Scenario A (pseudonymous to Tenant), Scenario B (identified to Tenant), or Scenario C (platform-level, handled separately). This determines the depth of review pathway.
- Layer 2 — Core checks (Section 4). Fifteen checks applied to every questionnaire.
- Layer 3 — Triggered modules (Section 5). Additional checks that fire when content meets a trigger condition.

Tenant-type lens (Section 6) is applied throughout — six contextual prompt sets per tenant category.

Delivery modes

Delivery modes (added v1.2.0): the Framework applies across WhatsApp delivery (push_template, inbound_menu, inbound_standalone, menu_subquestionnaire) and Wakamoso mobile-app delivery (mobile_app_tenant_questionnaire and mobile_app_service). The notification surface (Surface 3) is render-agnostic across both channels. Persistent mobile-app services route under the Tenant Zero pattern, with Wakamoso as responsible party and one EAR per service version.

3. Scenario triage

The first decision in every assessment is the scenario type. Scenario determines pathway, turnaround, and depth.

3.1 Scenario A — Pseudonymous to Tenant

- Audience: Wakamoso user base via demographic targeting, wa.me trigger arrivals, or Tenant Zero targeting on anonymised criteria.
- What Tenant sees: pseudonymous response stream only, no identifying information.
- Pathway: Light-Touch (turnaround: 24 hours).
- Core checks: C1, C2, C3, C4, C8, C11, C12 only. Other Core checks omitted unless content triggers a module.

3.2 Scenario B — Identified to Tenant

- Audience: Tenant-uploaded list (fixed schema: first name, surname, mobile, optional email), OR questionnaire contains PI-collecting questions, OR event-captured attendees followed up by the Tenant.
- What Tenant sees: response stream linked to identity (the Tenant already had the contact information).
- Pathway: Standard (turnaround: 7 Business Days).
- All fifteen Core checks plus any triggered modules apply.

3.3 Scenario C — Platform-level

- Audience: Wakamoso platform users intentionally engaging with platform features (Wallet, profile, CV).
- Pathway: handled at platform engineering level, not per-questionnaire. The Wakamoso POPIA Opt-In governs.
- Not in scope of this Framework's per-questionnaire review.

4. Core checks

Fifteen checks applied to each questionnaire (subject to Scenario A's reduced subset). In the ordinary course, any red finding blocks fielding and any amber finding requires documented mitigation.

C	Check	POPIA / ethics basis	Pass criterion
C 1	Purpose specification	S13 — specific, explicitly defined, lawful	Questionnaire purpose is named, lawful, consistent with Tenant's registered category.

C	Check	POPIA / ethics basis	Pass criterion
C 2	Lawful basis	S11 — consent / contract / legitimate interest / public task	Tenant identifies the Section 11 ground. For legitimate interest, an LIA is on file. For consent, the POPIA Opt-In covers the purpose.
C 3	Minimality and Personal-flag accuracy	S13(1)(c) — adequate, relevant, not excessive	Every question serves the stated purpose. Every field that collects Personal Information is marked 'Personal'.
C 4	Notification at collection	S18 — data subject informed	POPIA s.18 notification for this submission is delivered by the first paragraph of the questionnaire body (Surface 3); for mobile-app delivery this is the questionnaire or service intro screen - the same render-agnostic surface. It names the responsible party (the Tenant, or Wakamoso for Tenant Zero and mobile-app services) and states the purpose. The platform POPIA Opt-In is a separate canonical surface affirmed once per respondent and is not edited per submission.
C 5	POPIA Opt-In alignment	POPIA Opt-In v2.0 scope	Processing falls within the scope of the platform POPIA Opt-In v2.0. Where the Tenant's registered category matches a recognised lens and the declared purpose is consistent with that lens, the check passes without further scope analysis; persistent Wakamoso mobile-app services (Tenant Zero) pass by construction. It escalates to human judgement only where the category is unrecognised or the declared purpose is incompatible with the registered lens.
C 6	Special-category screening	S26 — special personal information (closed enumerated list)	Each question screened against the S26 enumerated attribute list (religion, race or ethnic origin, trade union, political persuasion, health, sex life, biometric, criminal behaviour). Trigger fires the Special-Category Module only where an enumerated attribute is DIRECTLY collected. Aggregation of non-special data into a sensitive socio-economic profile is OUT OF SCOPE for this check — it routes to C14 (cumulative profiling) and C15 (lifecycle) with an enhanced_safeguards condition to Deputy IO review, not a regulator_application condition. Section 27 authorisation is engaged only where actual Section 26 attributes are processed.
C 7	Direct marketing classification	S69 — separate consent	Genuine research passes. Marketing-dressed-as-research triggers stricter lawful basis.

C	Check	POPIA / ethics basis	Pass criterion
C 8	Reward proportionality and framing	Coercion check; compensation framing	Reward is proportionate. Reward is framed as compensation for participation, not payment for data.
C 9	Re-identification and location precision	Data Publication Policy floors	Regime-aware under v1.2.1 (DEC-2026-002). Each location-bearing field carries a declared regime (personal, public_listing, evidentiary, service_delivery, operational_asset, engagement_mapping). Personal and engagement-mapping data follow the platform-controls path - the matching engine reads full precision while visualisations and exports apply 2-decimal precision plus jitter and a $k \geq 5$ floor, with account-deletion purge; public-listing and operational-asset data publish at full precision by design; evidentiary and service-delivery data are disclosed only to named lawful recipients and not published. Where no regime is declared, every location field defaults to personal at 2-decimal precision. Aggregate publication respects the $k \geq 5$ / $k \geq 10$ floors. The check confirms the declared regime matches the question and the protective path is followed; it does not require truncation of the captured value.
C 10	Respondent welfare	Harm avoidance, non-stigmatisation	Questionnaire does not re-traumatise or stigmatise. Crisis content triggers welfare module and referral info.
C 11	Methodological integrity	Sanity check	No leading questions, double-barrelled questions, biased framing.
C 12	Disclosure to data subject	Plain English	Intro and per-question text accessible to target audience reading level.
C 13	Trigger Word appropriateness (where applicable)	Trigger Registry governance	Applies to WhatsApp delivery only; for mobile-app-only delivery it is not applicable. For WhatsApp delivery, the single inbound trigger word is unique, brand-appropriate, non-offensive, not a third-party trademark, and not a reserved keyword (STOP, START, HELP, UNSUBSCRIBE); the push-template name, body, and image are captured at intake and approved platform-side at send time.
C14	Cumulative profiling assessment	S13 (purpose specification) read with S14 (further processing compatibility)	For submissions targeting opted-in users with prior submissions on file (Tenant Zero in particular), the cumulative cross-submission profile is assessed. The check asks: does the cumulative profile across this and prior

C	Check	POPIA / ethics basis	Pass criterion
			submissions raise a separate processing-purpose question under S14 that is not already covered by the platform POPIA Opt-In? Where the cumulative profile creates risks not addressed by the current Opt-In scope (e.g. longitudinal economic vulnerability mapping, behavioural classification, persistent socio-economic identity), the Tenant must propose either a re-consent flow OR a documented purpose-compatibility analysis under S14. Output condition is enhanced_safeguards routed to Deputy IO, not regulator_application.
C15	Retention, secondary use, and lifecycle	S13 (purpose), S14 (further processing), S18 (notification of retention)	Tenant has explicitly named in the submission package: (a) retention period for collected data; (b) downstream analytics use (descriptive, inferential, predictive); (c) AI model training implications, if any (including translation-quality testing payloads); (d) publication governance pathway (which datasets are published, at what aggregation floor, subject to which review); (e) deletion pathway (when and how data is purged). Each must be consistent with the stated primary purpose under S13 and the further-processing rules under S14. Silence on any lifecycle element is a REQUIRES_HUMAN_JUDGEMENT result with a documentation_fix advice item, not a fail.

5. Triggered modules

5.1 Special-Category Module

Fires when any question DIRECTLY collects a Section 26 enumerated attribute (religion, race or ethnic origin, trade-union membership, political persuasion, health, sex life, biometric information, criminal behaviour). Adds: Section 27 authorisation pathway assessment; Section 57 prior-authorisation consideration; per-attribute lawful-basis ground. Does NOT fire where non-special data aggregates into a sensitive socio-economic profile — that risk is handled under C14 (cumulative profiling) and C15 (lifecycle) with an enhanced_safeguards condition routed to Deputy IO review, NOT a regulator_application condition. The 'health information' framing for gambling, addictive behaviour, financial distress, food insecurity, or survival-time questions does NOT trigger this module unless the questionnaire collects an enumerated health attribute (medical diagnosis, treatment, prescription, or biometric).

5.2 Vulnerable-Population Module

Fires when target population includes minors, GBV survivors, people in crisis, food-insecure households, refugees, or any cohort the Tenant flags as vulnerable, OR when question content surfaces vulnerability. Adds: consent capability check (competent-person consent for minors); reframed question wording; embedded referral; tightened minimality; $k \geq 10$ publication floor for any cell involving the vulnerable cohort.

5.3 C3B Framework Alignment Module

Fires when the questionnaire claims alignment to the C3B Community Capability Baseline, the Simmering Pot, or the SDG Alignment Rubric. Adds: domain and question-level mapping verification; scoring integrity check against canonical rubric v14; Tier One survival flag handling protocol; domain-coverage claim verification; correct naming ('C3B-derived' vs 'C3B Baseline').

5.4 Cross-Border Module

Fires when data will leave South Africa. Adds: Section 72 ground identification; Tenant's recipient DPA confirmation; data-subject notice of cross-border transfer in the questionnaire intro.

5.5 Cohort-Size Floor

Independent of modules above. Applies on every fielding regardless of Scenario.

- Cohort size < 50 : dispatch blocked. Tenant must broaden the criteria.
- Cohort size 50-200: dispatched, with publication floor automatically raised to $k \geq 10$ regardless of attribute type.
- Cohort size > 200 : standard $k \geq 5$ / $k \geq 10$ floor based on attribute type.

6. Advice to Tenant

Every EAR includes an Advice to Tenant section that pairs each finding with one or more concrete remediation steps the Tenant can take to improve the EAR outcome. Advice is constructive, specific, and prioritised. The intent of this section is to convert governance findings from a list of failures into a path the Tenant can walk to compliance.

Three categories of advice:

- `tactical_questionnaire_fix` — a direct change to the questionnaire content (remove a question, rephrase a leading question, add a welfare signpost screen, insert a screening gate, replace a precise field with a coarse one). Tactical fixes are the lowest-cost path to remediation.
- `structural_methodology_fix` — a change to scope, sampling, or methodology (split a dual-purpose survey into two, narrow the target cohort, change delivery channel). Structural fixes are higher cost but sometimes the only path forward.
- `documentation_fix` — a change to what the submission package documents, without touching the questionnaire (state the retention period, state the lawful basis explicitly, attach an LIA, name the downstream analytics use). Documentation fixes are zero-impact on the questionnaire but often close `REQUIRES_HUMAN_JUDGEMENT` findings to `PASS`.

Each advice item is paired with an 'if-actioned' projection: which finding it addresses, which condition it would close, and how it would change the recommended verdict. The Tenant should

be able to read the Advice section and produce a revised submission package in a single editing pass.

Where a finding stems from documentation silence (e.g. 'the submission does not name the data retention period'), the advice clearly states that the gap can be closed by adding the required information rather than changing the questionnaire. Where a finding stems from a design choice (e.g. 'Q2 collects precise GPS'), the advice states the specific change required and any acceptable alternatives.

Tone: constructive, not punitive. The Office of Ethics Review exists to help Tenants ship compliant research, not to obstruct legitimate work. Advice should make this orientation visible.

7. Tenant-type lens

Each of the six tenant categories carries contextual prompts that direct attention to category-specific risks. The lens does not generate additional pass/fail findings; it informs the depth of attention paid to specific Core checks.

- Business / private sector — marketing-research overlap; workforce power-asymmetry; competitor probing limits.
- Loyalty programmes — member-privacy expectations; first-party data scope; reward in loyalty currency.
- Research agencies — onward disclosure to end-client; methodological standards; co-branding restrictions.
- Academia — IRB equivalence; Section 32 safeguards; Data Access Agreement requirement.
- Government / municipalities — political neutrality; public-task lawful basis; retaliation risk.
- Development / NGOs — vulnerable-population presumption; funder-driven question patterns; programme-evaluation reflexivity.

8. Decision logic

Profile	Decision
Scenario A, all Core greens, no modules fire	APPROVED via Light-Touch. EAR filed. Tenant fields within 7 days.
Scenario B, all Core greens, modules fire green	APPROVED via Standard pathway. EAR filed.
Any amber, no red	APPROVED-WITH-CONDITIONS. Tenant addresses each amber; Office confirms; fielding proceeds.
Single red (other than POPIA Opt-In misalignment)	RETURNED-FOR-REVISION.

Profile	Decision
POPIA Opt-In misalignment red	DECLINED unless Tenant reframes within Opt-In scope.
Multiple reds, or red on welfare	DECLINED. Tenant invited to reconceive.

9. AI-assisted assessment

Each assessment is initiated through the Ethics Assessor Coworker skill (specification v1.0). The skill produces a structured first-pass assessment against the Framework. The Deputy Information Officer reviews, modifies or overrides as needed, and signs the binding decision. The AI is the analyst; the human is the decision-maker. AI-assisted review is calibrated against the published Framework but is not neutral: it carries the interpretive, prompting, training, and framing bias inherent in any model-driven process. Calibration is reviewed by the Advisory Panel at each sitting.

10. Sign-off

Approved by:

Mélani Prinsloo — Information Officer (CEO)

Date: _____

Michael Matthews — Deputy Information Officer (CIO)

Date: _____

Reviewed by internal counsel:



Malusi Mazibuko — Internal Counsel, Wakamoso Africa (Pty) Ltd

Date: 08/06/2026

Master Adoption

This page records Wakamoso Africa (Pty) Ltd's adoption of the master version of the Wakamoso Ethics Framework v1.2.1. From the adoption date below, this is the canonical master text. Substantive changes require a new version number and a new Master Adoption record.

Document: Wakamoso Ethics Framework

Version: v1.2.1

Adoption date: 27 May 2026

Responsible Party: Wakamoso Africa (Pty) Ltd, registration 2022/762282/07

Information Officer: Mélanie Prinsloo, Chief Executive Officer

Deputy Information Officer: Michael Matthews, Chief Information Officer

Internal Counsel: Malusi Mazibuko, Internal Counsel

Information Regulator Registration Number: 2026-018959

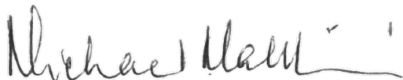
Adopted by Wakamoso Africa (Pty) Ltd:



Mélanie Prinsloo — Information Officer (CEO)

Date: 08 June 2026

Witnessed by the Deputy Information Officer:



Michael Matthews — Deputy Information Officer (CIO)

Date: 08 June 2026

Reviewed by internal counsel:



Malusi Mazibuko — Internal Counsel

Date: 08/06/2026

Once adopted, the fully countersigned document is uploaded to the Governance Dashboard and recorded as the adopted master version. Subsequent changes require a new version and a new Master Adoption record.