

# Wakamoso Ethics Framework v1.2.1

*A plain-language reference*

Governance / POPIA Compliance | 8 June 2026 | Reference only - not the binding instrument

This is a readable guide to how the Wakamoso Ethics Framework reviews questionnaires. The binding version is the formal Ethics Framework v1.2.1 document; the operative rubric the software runs is maintained in the Governance Dashboard. This guide explains, in plain terms, what the review actually checks and why.

## How a review works

Every questionnaire a Tenant wants to field is reviewed before it reaches any respondent. The review is human-led and AI-assisted: an AI "Ethics Assessor" produces a structured first-pass assessment against this Framework, and the Deputy Information Officer then reviews it, adjusts or overrides it where needed (with reasons), and signs the binding decision. The AI is the analyst; a person is always the decision-maker. Routine submissions are turned around within 7 business days.

Findings come in two kinds. POPIA findings (law) bind the decision and can attach conditions. Ethics findings (good practice) are advisory. Each review ends in one of five outcomes:

- APPROVED - cleared to field.
- APPROVED WITH CONDITIONS - cleared once the Tenant meets stated conditions.
- RETURNED FOR REVISION - changes needed; resubmit.
- DECLINED - not permitted on the Platform (the Tenant may appeal).
- APPROVED PENDING REGULATOR - cleared but cannot field until the Information Regulator responds (Section 57 cases).

## Two delivery channels, three surfaces

Questionnaires reach people two ways: over WhatsApp, or in the Wakamoso mobile app. The rules are the same in both; only the screens differ. Every respondent journey has three text "surfaces":

- Surface 1 - the entry path (a WhatsApp push message or trigger word, or an app tile / menu).
- Surface 2 - the platform POPIA Opt-In, agreed once per person at first contact. There is no separate consent form for each questionnaire.
- Surface 3 - the questionnaire (or service) intro screen. This is where the "who and why" notice lives, and it is what the notification check looks at.

## The three layers of the review

- Scenario triage - first, what kind of fielding is this? Scenario A (responses are pseudonymous to the Tenant) gets a light-touch review of a reduced check set. Scenario B (the Tenant already knows the people, e.g. an uploaded contact list) gets the full review. Scenario C (platform-level features) is handled at engineering level, not per questionnaire.
- Core checks - fifteen checks (C1-C15) applied to each questionnaire. These are the heart of the review and are listed below.
- Triggered modules - extra checks that switch on only when the content calls for them.

A tenant-type lens runs throughout, tuning attention to the kind of organisation (business, loyalty programme, research agency, academia, government, or NGO).

### The fifteen core checks, in plain English

| #  | What it asks   | Type  |
|----|--|-------|
| C1 | Is the purpose clearly stated, lawful, and consistent with what this Tenant does?  | POPIA |
| C2 | Is there a valid legal ground to process the data (the platform Opt-In, a contract, legitimate interest with an assessment on file, or a public task)?   | POPIA |
| C3 | Does every question serve the purpose (nothing excessive), and is every personal field correctly flagged "Personal"? Standard demographics (age band, gender, province) need no special justification. | POPIA |
| C4 | Does the questionnaire intro screen (Surface 3) say who is collecting the data and why?  | POPIA |
| C5 | Does the intended use fall within the Opt-In people already agreed to?<br>Recognised tenant types and Wakamoso's own services pass automatically.  | POPIA |
| C6 | Does any question directly collect a Section 26 special category (religion, race or ethnicity, union membership, politics, health, sex life, biometrics, criminal behaviour)? Combining                | POPIA |

|            |  |        |
|------------|--|--------|
|            | ordinary questions does not make them special.   |        |
| <b>C7</b>  | Is this genuine research, or marketing dressed up as research? Marketing needs separate consent (Section 69).  | POPIA  |
| <b>C8</b>  | Is any participant reward proportionate and framed as compensation for time - not payment for personal data?   | POPIA  |
| <b>C9</b>  | Is location data handled according to its purpose ("regime")? Personal locations are protected (about 1 km precision, jitter, and a small-group floor); public listings show full precision by design. | POPIA  |
| <b>C10</b> | Could the questions harm, re-traumatise, or stigmatise respondents? Sensitive topics trigger welfare safeguards.   | Ethics |
| <b>C11</b> | A sanity check on method: no leading, loaded, or double-barrelled questions.   | Ethics |
| <b>C12</b> | Is the language plain enough for the intended audience to understand?  | POPIA  |
| <b>C13</b> | For WhatsApp only: is the trigger word appropriate and not a reserved keyword? Not applicable to mobile-app-only delivery.   | POPIA  |
| <b>C14</b> | Across repeated questionnaires to the same people over time, does the accumulation build a profile beyond the stated purpose?  | POPIA  |
| <b>C15</b> | Has the Tenant said how long data is kept, what it may be reused for, and how it is disposed of?   | POPIA  |

*POPIA checks bind the decision; the two Ethics checks (C10 welfare, C11 method) are advisory.*

## Triggered modules (switch on by content)

- Special-Category Module - fires when a question directly collects a Section 26 attribute; adds Section 27 / 57 requirements.
- Vulnerable-Population Module - fires when the audience includes minors, GBV survivors, people in crisis, and similar; adds extra safeguards.
- Cross-Border Module - fires when data will leave South Africa; adds Section 72 requirements.
- A cohort-size floor also applies to every fielding (very small groups are blocked or have publication floors raised) to prevent re-identification.

## Location data, explained (C9)

Rather than blunt-truncating every coordinate, v1.2.1 handles location by what it is for. Each location question is tagged with a "regime":

- Personal / engagement-mapping - protected: the matching engine may read full precision, but anything shown to staff, shared with Tenants, or published is reduced to about 1 km precision with jitter and a small-group ( $k \geq 5$ ) floor; deleted when the account is deleted.
- Public listing / operational asset - shown at full precision on purpose (that is the point of a listing or asset registry).
- Evidentiary - kept precise for evidence, never aggregated, shared only with specific lawful recipients.
- Service delivery - precise enough to deliver the service, shared only with the named provider, not published.

Where a questionnaire says nothing, the most protective setting (personal, about 1 km) is assumed. Wakamoso does not do background or continuous location tracking.

## A few principles worth remembering

- One consent: people affirm the POPIA Opt-In once - there is no per-questionnaire consent form.
- Section 26 is a closed list: combining ordinary data does not turn it into special-category data (but cumulative profiling is watched under C14).
- Same rules, two channels: WhatsApp and the mobile app are assessed identically; only the screens differ.
- Help, not gatekeeping: every assessment pairs findings with practical advice on how to fix them.

## What changed since v1.1 (brief)

v1.2.1 reflects calibrations through the location-regime decision (DEC-2026-002): C14 and C15 became core checks; C5 gained recognised-category and Wakamoso-service auto-passes; C4 was clarified to the render-agnostic Surface 3; mobile-app delivery modes were added (v1.2.0); and

C9 became regime-aware (v1.2.1). The full calibration history is in the formal Framework v1.2.1 document.