



PRESERVING LIFE: COMPANION COACHING ASSISTANT GUARDRAILS AND ATTESTATION

This document sets out safety, ethics, and transparency guardrails for the Companion Coaching Assistant. It clarifies what the Companion will and will not do, how it avoids impersonation, how it escalates when risk is detected, and how professional coaching standards inform tone and boundaries. It supports diligence and does not amend any agreement. If there is a conflict, the EULA and any signed Order Form or Data Processing Addendum (DPA) prevail. Use this document to understand the always and never rules, the stop-then-support escalation pattern, privacy and data minimisation measures, and the division of responsibilities when a general-purpose model is configured.

1. PURPOSE AND ETHOS

Superworker's Companion Coaching Assistant ("the Companion") is designed to move work forward while protecting dignity, psychological safety, and life. This document describes the non-promissory guardrails that govern the Companion's behaviour. It supports due diligence and does not amend any agreement. If this document conflicts with your EULA or Order Form, those documents prevail.

Our stance is simple: do no harm, be transparent, keep humans in control, and escalate to real people when safety requires it.

2. TRANSPARENCY AND NON-IMPERSONATION (EU AI ACT ALIGNMENT)

Modern AI can produce responses that read like a human. To prevent confusion and in line with the EU AI Act transparency duties, the Companion will:

- Clearly identify itself as an AI assistant at the outset and in settings.
- Avoid pretending to be a human, a named individual, or a specific professional.
- Disclose when content is AI generated, or AI assisted where such disclosure is relevant to user understanding.
- Preserve human oversight. Final decisions remain with people.

We do not design for prohibited practices such as manipulative techniques that cause harm, social scoring, or unlawful biometric categorisation. Where enterprise-specific policies require additional disclosures, we support those configurations.



info@superworker.co



410 Lynnwood Road,
Lynnwood,
Pretoria,
0081



3. PROFESSIONAL COACHING STANDARDS (ICF / GLOBAL CODE / COMENSA)

We align our coaching behaviours with established professional frameworks and will not present the Companion as a substitute for a qualified coach or clinician.

- International Coaching Federation (ICF) Code of Ethics: contracting, confidentiality, clear boundaries, and referral obligations (<https://coachingfederation.org/ethics>).
- Global Code of Ethics (Association for Coaching and EMCC Global): enforceable standards and complaints procedures (<https://www.globalcodeofethics.org/>).
- COMENSA Code of Ethics (South Africa): where relevant to South African deployments (<https://www.comensa.org.za/>).

These frameworks shape tone, boundaries, and escalation behaviours. This document does not claim accreditation.

4. SCOPE BOUNDARIES: WHAT THE COMPANION WILL AND WILL NOT DO

WILL

- Offer assistive coaching prompts that help users make progress on work and learning tasks.
- Provide options ("Show Me" for journeys, "Ask Me" for coaching steps) with micro-commitments and clear next actions.
- Use bias-aware nudges, for example planning-fallacy buffers, to reduce rework and stress.
- Surface policy-relevant reminders (privacy, compliance, safety) in context.

WILL NOT

- Impersonate a human, provide therapy, or give medical, legal, or financial advice.
- Encourage unsafe activities or bypass enterprise policies.
- Collect sensitive personal information unless expressly required and authorised by the customer.
- Continue coaching when safety concerns arise. It pauses and escalates.

5. SAFETY AND ESCALATION: PRESERVING LIFE

When the Companion detects distress or risk, it follows a conservative "stop, then support" pattern:

- Pause coaching immediately and acknowledge the user's state in clear, kind language.
- Run a short, non-diagnostic safety check ("Are you safe right now?").
- Provide warm referral options, for example EAP, local crisis lines, or designated organisational contacts, and offer to stay connected while the user reaches out.
- Log minimal details required for audit, store securely, and share only on a need-to-know basis (POPIA and GDPR principles).
- Notify the customer's designated safeguarding contact per configuration.
- Resume only when appropriate and with user consent, otherwise remain paused.

The escalation posture is informed by ICF ethics (referral duties), WHO mhGAP, and MHFA guidance. We do not diagnose or treat.



6. PRIVACY AND DATA MINIMISATION

Superworker is built with privacy in mind. Customers act as Controllers or Responsible Parties; Superworker acts as Processor or Operator. We minimise personal data in prompts and logs, and we encourage customers to avoid special-category data unless strictly necessary and authorised.

- Hosting on Microsoft Azure in customer-selected regions. Encryption in transit and at rest.
- Minimal logging for reliability and abuse detection. No sale of personal data.
- No training of public models on customer inputs or outputs. LLM provider behaviour follows that provider's terms and tenant configuration.
- Anonymisation and pseudonymisation: redact direct identifiers in free text where feasible; use tenant-local pseudonyms when correlation is required.
- Role-based access, SSO and MFA support, and auditable administration.

7. AI MODEL USE (GPAI AND PROVIDER RESPONSIBILITIES)

Where a general-purpose model, for example a GPT-class model, is configured by the customer, that model provider is independently responsible for its obligations under applicable law. Superworker documents the configured provider and selects available options that restrict training on customer data and limit retention where such options exist.

We scope retrieved context to authorised sources and bounded windows, and we transmit requests over TLS.

8. GUARDRAILS: ALWAYS / NEVER

ALWAYS

- Be transparent about being an AI system.
- Offer a human-in-the-loop path and clear exits to people-run channels.
- Keep interventions small, safe, and measurable. Celebrate progress without shaming.
- Escalate early when safety risk is suspected. Log minimally and protect confidentiality.

NEVER

- Pretend to be human or a specific named person.
- Pressure users into disclosure of personal or sensitive information.
- Provide therapy or regulated professional advice.
- Bypass enterprise policies or encourage unsafe behaviour.

9. CONFIGURATION (PER-TENANT CONTROLS)

- Safeguarding contacts and crisis lines for warm referrals.
- Preferred disclosures and banners for transparency.
- Local ethics frameworks (ICF, Global Code, COMENSA) as applicable.
- Retention windows and access roles aligned to the customer policy.



10. REFERENCES AND CONTACT

- International Coaching Federation — Code of Ethics: <https://coachingfederation.org/ethics>
- Global Code of Ethics (AC and EMCC): <https://www.globalcodeofethics.org/>
- WHO mhGAP Intervention resources: <https://www.who.int/teams/mental-health-and-substance-use>
- Superworker website: <https://www.superworker.co>
- General enquiries and safety escalation contact: info@superworker.co