

POLICY

Policy on the use of artificial intelligence

REFERENCE	REVISION	DATE	APPROVED
RCL-P-2026-001	A	16 October 2026	Renier Cloete

1 Purpose

This policy says what we use artificial intelligence (AI) for, how we keep its output safe to rely on, and what happens to our clients' information. We give it to every client and to anyone else who relies on our work. It follows the Institution of Structural Engineers' AI guidance for organisations and its guidance on large language models (LLMs) in calculations. We do not seek ISO/IEC 42001 certification; we would reconsider if a client asked for it.

2 What we use AI for

We use Anthropic's Claude models, through Claude Code and claude.ai. We use AI in three ways.

Inside RCL. To write and test our software, to find clauses in codes of practice and textbooks in our reference library, and to help with research and drafting.

In work others rely on. Our reports, articles and calculation software are made with AI assistance. In all of it:

- An engineer at RCL is responsible for every output, whatever tool produced it.
- Calculation results come from our own deterministic software, which is checked against its published verification manual. An LLM writes and explains code; it does not produce the numbers in a calculation.

- Work produced with an LLM is checked as a graduate's work would be: the method, the references and the results. We do not take an LLM's statement that a task is complete as evidence that it is.
- An LLM is never the independent checker. Where a check must be independent, it uses different software, and the LLM used in the design is named in the checking record.
- Each published piece says that it was written with AI assistance.

In tools we provide to others. We have not yet released an AI tool to clients or other engineers. Before we do, we will test it on a published benchmark, document what it cannot do, name the model it is built on, and tell its users that its output is AI-generated and must be checked.

3 Your information

- We handle personal information in line with the Protection of Personal Information Act, 2013 (POPIA).
- We put a client's project information into an AI tool only with that client's written consent, given for the project, and only under Anthropic's commercial terms (the Claude API), under which Anthropic does not train models on what we send and keeps it for 30 days.
- Before project material goes into an AI tool, we remove the names and contact details of people that the work does not need.
- In client work we do not send feedback reports or share session transcripts with Anthropic, since those are kept for longer.
- Our other work, on our own and public material, uses Claude accounts on Anthropic's consumer terms with the setting that allows use for training switched off; on those terms Anthropic also keeps conversations for 30 days.
- We do not use AI tools on projects under a non-disclosure agreement, or on projects whose layout could help someone intending harm (critical infrastructure, transport structures, government, financial, health and power buildings), unless the client asks us to in writing.
- We do not upload codes of practice, textbooks or other material we do not own to train or extend an AI model.

- We keep project information in project folders with access limited to the people working on the project. When the project ends we delete it from AI tools: conversations through Claude's deletion controls, and session records from our own machines.

4 Keeping track of AI tools

- For each job and each tool we record the model and version used.
- We re-test our tools when the underlying model changes, and at least once a year: a result obtained with one model is not evidence about another.
- A tool designed to one generation of a code of practice is retired or revised when that code is replaced.

5 Responsibility and concerns

Renier Cloete is RCL's AI lead and is responsible for this policy. Anyone with a concern about our use of AI, including a concern about structural safety, can raise it with him at renier@rcl.co.za. A tool under a safety concern is withdrawn from use until the concern is resolved, and the work it touched is reviewed.

6 Review

We review this policy every year, and sooner when the tools, the law or the guidance change.

Annex: risks and how we manage them

RISK	HOW WE MANAGE IT
An LLM makes an error in code, a calculation method or a reference	Calculations come from verified software; LLM work is checked as a graduate's; references are checked against the code itself
An AI agent reports work as complete when it is not	Completion is judged by checks and review, never by the agent's own report
Designer and checker share the same blind spots	Independent checks use different software; the model used in design is recorded
Client information leaves our control	Written consent per project; names and contact details not needed removed first; client work only under commercial terms (no training, 30-day retention); no feedback or transcript sharing; sensitive projects excluded
Material we do not own is passed to an AI model	Reference material is used for lookup only; nothing is uploaded to train a model
A model update changes a tool's behaviour	Model and version recorded; tools re-tested on each model change and yearly
A code of practice changes	Tools built to the old code are retired or revised