
How VerkTools Handles Your Data

Five commitments

- 1 • Your data is never training data.** We use commercial AI APIs configured so your information does not train the provider's models. We never use your data to build, tune, or improve anything for any other client.
- 2 • Your data lives alone.** You occupy an isolated cell — your own datastore, your own credentials, your own deployment. No shared databases, no commingling. This is enforced by architecture rather than by policy, which means it does not depend on anyone remembering it.
- 3 • Access is minimal and logged.** Named individuals only, least privilege, unique credentials. Every access to your data is attributable and reviewable.
- 4 • You can always leave with everything.** Your data and outputs export to standard formats — CSV, Excel, PDF — on request. On termination we return everything, then delete it, then certify the deletion.
- 5 • If something goes wrong, you hear it from us first.** Any suspected incident affecting your data: notification within 72 hours of confirmation, telling you what happened, what is affected, and what we are doing about it.

How AI is allowed to work here

We build with AI. These are the limits we put on it.

A named human approves every client-facing output. Cost math, tool outputs, calibration changes — each one carries a human approver by name. AI drafts; humans own. This accountability is not delegable to a model.

Tools flag, draft, and calculate. They do not decide. No system we build executes a business decision autonomously. The human-in-the-loop gate written into every statement of work is a governance control, not a product feature.

Every structured claim traces to its source. Model-generated content is always distinguishable from fact you supplied — you can tell, in any output, which is which.

Prompts containing your confidential data stay yours. They are scoped to your engagement, stored in your cell, and never reused across clients.

Business logic is separated from the model. Deterministic calculation is isolated from the AI layer. Where the system extracts a value it is not confident about, that field routes to a human rather than being guessed. Any material change of model or provider is logged against your engagement.

What we will not do. No client data in any public or consumer AI tool. No autonomous external communication on your behalf. No biometric, surveillance, or employee-monitoring builds.

Who else touches your data

We maintain a register of every service in the data path — provider, role, data-processing terms, training-use status, and region.

You can request that register at any time and we will send it.

Adding any new vendor to the data path requires a documented review against the five commitments above. A service that has not passed that review does not receive your data.