



Trust & Security Overview

Agentic Solutions, LLC | v1.1 – July 2026

Agentic Solutions builds and implements AI-powered solutions for businesses — which means clients trust us with their data, systems, and workflows. This sheet sets out, in plain language, the concrete measures behind that trust. We're glad to walk through any of them in detail.

Our Approach to AI

AI is our craft, and we hold it to professional standards. Every AI-assisted deliverable is reviewed by a member of our team before it reaches a client. Client data is only ever processed through AI providers with whom we hold commercial agreements and data processing terms — never through free or consumer-grade tools. And we're always transparent about where and how AI is used in our work.

How We Handle Your Data with AI

- **Anthropic is our primary AI provider.** Our core AI work runs on Anthropic's Claude under commercial terms that include a data processing agreement. Any additional AI providers we use are accessed only under commercial or enterprise terms — never free or consumer tiers — and open-source models, when used, run in our own controlled environment.
- **No training on client data.** Across all of our providers, we operate only under terms and configurations that exclude client data from model training.
- **Bounded retention.** Data we submit for AI processing is automatically deleted by the provider, with retention capped at 30 days — and we configure shorter retention windows wherever the selected provider supports them.
- **Backed by our agreements.** These commitments come from our commercial agreements with our providers, not just internal policy; relevant terms are available to clients on request.

Full policy: agentic solves.com/docs/ai-data-policy.pdf.

How We Secure Data on Our Side

- **Client environment isolation.** Our environments are designed to keep each client's data and working environment separate, so work for one client is not performed with access to another's.
- **Separate, named accounts.** No shared logins. Access is individual, attributable, and revoked promptly when a role changes, with team members limited to the data their current work requires.
- **Credential and device security.** Credentials are managed through a password manager, and all devices handling client data are encrypted.
- **Clean offboarding.** When an engagement ends, client data is returned or securely deleted per our agreement.
- **Confidentiality and incident notification.** We're glad to enter into an NDA at any stage, and we notify affected clients promptly of any security incident affecting their data, consistent with our agreement.

Full details: agentic solves.com/docs/security-posture.pdf.

Insurance Coverage

- **Professional Liability (Errors & Omissions):** \$1,000,000 per claim / \$1,000,000 aggregate
- **Cyber Liability:** \$1,000,000 per claim / \$1,000,000 aggregate
- **General Liability:** \$1,000,000 per occurrence / \$2,000,000 aggregate

Certificates of insurance are available on request.

Questions or due-diligence requests? Contact us at contact@agentic solves.com. The latest version of this document is available at agentic solves.com/docs/agentic-solutions-trust.pdf.

This overview is provided for information only and is current as of the date above. It does not modify or supersede the terms of any signed agreement, which govern our engagement. Details are subject to change.