How Risk Crew uses AI

A plain-English statement of our practice

AI is now part of how serious cybersecurity work gets done. We use it where it makes our testing and advisory work faster, deeper or more thorough, and we are open with our clients about where and how. This piece explains our practice in plain English.

Risk Crew

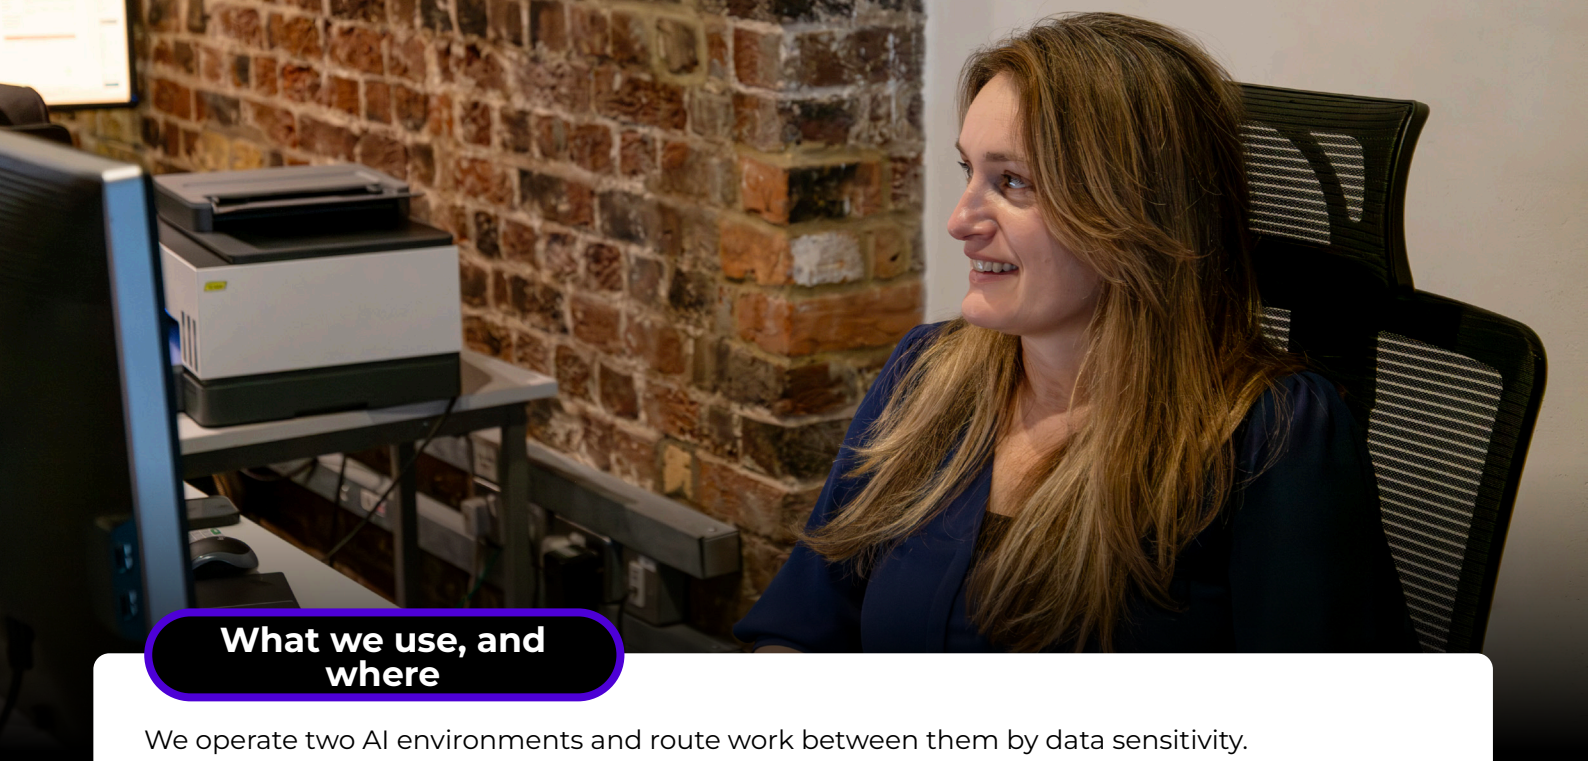

What we use, and where

We operate two AI environments and route work between them by data sensitivity.

External AI: Currently Anthropic's Claude, processed in the Anthropic EU region under a contract that excludes the use of our submitted content for model training. We use Claude for methodology research, general technical reference, review of payload patterns against published vulnerability classes, draft language for non-client-specific report sections, and internal productivity tasks.

In-house AI: Our own infrastructure, on-premise, fully controlled by Risk Crew. Nothing submitted to this environment leaves our network. All work involving live client testing data, reconnaissance output, application responses, proof-of-concept analysis, captured credentials, or any artefact produced from access to a customer's systems, is handled here, not in any external service.

Embedded AI: Some commercial security tools we use have AI capabilities built in and we think that will increase. We assess each tool's data handling against our standards before allowing it on client work.

What we do not use AI for

Regardless of environment, AI is never used by Risk Crew to perform any of the following:

- Final risk scoring or CVSS attribution of findings.
- Final attestation that a vulnerability exists or has been remediated.
- Authorship of the executive summary or any management-facing narrative without senior consultant authorship and sign-off.
- Approval or sign-off of any deliverable
- Submission of live exploit payloads, working proof-of-concept code, captured credentials, or unpatched vulnerability detail to any external AI service.

Every external deliverable carries a named Risk Crew tester and QA lead. AI cannot sign off, score, attest or approve. Professional responsibility for the work rests with the people who do it, not the tools they use.

Risk Crew

Your Rights as a Client

Right to information. On reasonable request, you may receive the engagement's AI tool register and a summary of how each tool was used.

Right to opt out. You may, in writing before commencement, require delivery without the use of any external AI service. We will accommodate such requests, and we reserve the right to discuss any commercial or timeline implications before confirming the engagement.

Right to additional safeguards. Where you have specific contractual or regulatory obligations affecting AI sub-processors, we will agree any additional controls before commencement.

Questions

We expect this practice to evolve as the technology, the standards and the regulatory landscape develop. We will keep this statement current and publish a dated version each time it changes. Questions or comments are welcome.

riskcrew.com

+44 (0)20 3653 1234

linkedin.com/risk-crew

5 Maltings Place,
169 Tower Bridge Road,
London,
SE1 3JB

Contact Risk Crew today and discover how AI Security can be safely implemented in your organisation.

Risk Crew