

How does Cyber Crucible reduce third-party and vendor risk for regulated buyers?

Short answer: By removing the risk rather than only managing it. The three facts that answer the hardest vendor-risk questions are: Cyber Crucible collects no customer content, credentials, or keys; protection runs locally and survives any backend outage; and its AI lives only in development, never on your endpoint. Together these shrink the surface a due-diligence review is designed to probe.

The three structural risk reducers

- **Zero content.** No customer files, credentials, or keys are ever collected. That single fact answers the PCI, confidential-data, and "what if you're breached" lines of questioning — the data is not there to lose.
- **Backend-independent protection.** Detect-Decide-Respond runs locally, with an effective recovery-time objective of zero on the endpoint. A total backend outage does not reduce protection.
- **Deterministic AI.** AI is used in development only; deterministic heuristics run at runtime. No live model, no drift, testable outcomes.

Independent reinforcement

Where independent attestation carries weight, Cyber Crucible points to what is real: Microsoft WHCP driver signing, and — for the managed backend — the third-party attestations of the infrastructure providers it relies on. With no SOC 2 of its own, these carry the load honestly rather than being oversold.

Revision #2

Created 2026-07-23 15:18:16 UTC by Dennis Underwood

Updated 2026-07-23 18:19:47 UTC by Dennis Underwood