

TVCRpro — Security Posture Brief

For licensee decision-makers

DOCUMENT	TVCR-SEC-001-B
VERSION	1.0
BUILD	Site v31.1 · Worker 1.0.0-phase3.14
DATE	May 5, 2026
CLASSIFICATION	Public
OPERATOR	Atomic Armstrong (sole proprietor, Ontario, Canada)
PATENT	USPTO Application 64/045,951 · Confirmation #7707 · Filed April 21, 2026
COMPANION	Technical brief for your security reviewer: TVCR-SEC-001-A

1. Why this brief exists

You are evaluating whether your team can rely on TVCRpro for measuring AI work, and whether the way it is built is something you can put in front of your own security and procurement people without surprises later. This brief gives you a plain-language read on the security posture so you can decide whether to advance the conversation. A companion technical brief (TVCR-SEC-001-A) is available for your information security reviewer.

2. What you are actually licensing

A measurement framework. Your team or your customer pastes a prompt and an AI response into a web form. The system returns an eleven-dimension score, a single composite ratio, and a coaching anchor. The eleven-dimension rubric is published openly on the website. The composite weighting and the token-normalization layer that turn those eleven scores into one ratio are protected and patent-pending. That separation — public rubric, protected math — is the core of the offering.

3. How data moves, in plain terms

When someone uses the analyzer:

The browser sends the pasted prompt and response to a single Cloudflare endpoint we operate at api.tvcrpro.com. That endpoint forwards the text to Anthropic's Claude API to score the eleven dimensions. The result comes back, our endpoint applies the protected weighting, and the final score is returned to the browser. The browser displays the score. Nothing is stored.

The website itself does not set cookies, does not run analytics pixels, and does not have an application database. There is no operator-side log of pasted prompts or responses. Cloudflare and Anthropic each see the traffic that flows through them under their published commercial terms, and Anthropic does not use API inputs or outputs to train its models.¹

4. What this means for your team

You can treat TVCRpro as a thin, transparent layer in front of Anthropic for the rubric step, plus a small piece of operator-protected math for the composite. There is no hidden third party. There is no analytics vendor in the chain. There are five named sub-processors and they are listed on the public security page with their trust-center links. If your security team has already approved Cloudflare and Anthropic for other workloads, most of the chain is already familiar to them.

The pasted text travels to Anthropic. It does not travel to any analytics, observability, or marketing vendor. There are no others.

5. The five sub-processors

Function	Provider	Where they sit
Hosting, DNS, edge security, the analyzer endpoint	Cloudflare, Inc.	U.S., with Canada and EU edge points
Operator email	Proton AG	Switzerland
Productivity and document workflow	Microsoft 365 (Canada tenant)	Canada
PDF rendering and signing	Adobe Inc.	U.S.
Inference for the eleven-dimension rubric	Anthropic, PBC	U.S.

The current list, with trust-center URLs, is published at tvcprpro.com/security. Any change to that list triggers a new published build of the website.

6. The protective controls already in place

The analyzer endpoint enforces a per-IP rate limit at Cloudflare's edge to deter abuse. The web pages enforce a strict content-security policy that prevents the analyzer from talking to anywhere except our own endpoint. Inputs are size-limited so oversized payloads cannot reach Anthropic. The Anthropic API key is stored as a Cloudflare Worker secret, never in source code, never in the browser bundle, and is rotatable. A monthly spend cap of USD 25 sits on the Anthropic account itself as a hard ceiling against runaway cost or misuse.

Every page of the site shows a build identifier in the footer, and a machine-readable build manifest is published at tvcprpro.com/build.json. That means the brief you are reading can be tied to a specific deployed version, and your reviewer can verify which version is live at the moment they look.

7. What this brief is not

This is the section your security reviewer will turn to first, so it is identical to the one in the technical brief.

- **Not a SOC 2 report.** No SOC 2 Type I or Type II audit has been performed. No third-party audit firm has been engaged. The operator is a sole proprietor; there is no audit committee.
- **Not a penetration test report.** No external penetration testing has been performed against tvcprpro.com or api.tvcprpro.com.

- **Not a formal vendor risk assessment.** Sub-processor selection was made by the operator; no third-party vendor risk management service has reviewed the chain.
- **Not a privacy program certification.** No GDPR Article 27 representative is appointed; no CCPA service-provider attestation is published; no standard data processing agreement template is published. Pilot engagements include a one-page DPA on request.
- **Not a continuity plan.** No documented incident response runbook beyond operator-direct notification through pilot@tvcrpro.com.

The trust artifacts that do exist today are: the patent filing, the architectural separation between the public rubric and the protected weighting, the public sub-processor disclosure, and direct access to the operator. These are intentional disclosures. The gaps above are real and the way to close them — for any specific licensee — is to scope a pilot that is sized to the trust artifacts in place rather than ones that are not.

8. What a pilot looks like, given the posture

A typical first engagement is small, time-bounded, and uses non-regulated data. The text your team pastes into the analyzer should not contain protected health information, payment card data, regulated personal information, or trade secrets you are not authorized to share with Anthropic. If your use case requires any of those, the right next step is a conversation about a private deployment path rather than the public analyzer at tvcrpro.com.

If your security team needs the SOC 2 conversation to begin before a pilot, that is fair, and the answer is that scoping for SOC 2 Type I has not yet been engaged. If your security team is comfortable with a scoped, non-regulated pilot in the meantime, the architecture, sub-processor list, and disclosures above are what is on offer.

9. How to verify everything in this brief without contacting us

Each of these is one command or one URL.

What you want to confirm	How
The patent application exists	Search Application 64/045,951 at patentcenter.uspto.gov
The site is on Cloudflare	<code>curl -I https://tvcrpro.com</code> and read the server header
The analyzer endpoint is alive	<code>curl https://api.tvcrpro.com/healthz</code>
The build version matches this brief	<code>curl https://tvcrpro.com/build.json</code>
The sub-processor list and identities	tvcrpro.com/security
The security headers on the site	<code>curl -I https://tvcrpro.com/analyzer</code>

If any of those return something different from what is described in this brief, the brief is the document that should be corrected.

10. How to move forward

- **Pilot scoping:** pilot@tvcrpro.com
- **Security review questions:** security@tvcrpro.com
- **Operator directly:** Atomic Armstrong — atomicarmstrong@proton.me

The companion technical brief (TVCR-SEC-001-A) goes deeper into the data flow, the worker controls, and the independent verification path; it is intended for your information security reviewer.

¹ <https://www.anthropic.com/legal/commercial-terms>

This brief is updated whenever the build identifier changes. The current build identifier is shown in the footer of every page on tvcrpro.com and at <https://tvcrpro.com/build.json>.

Next planned brief revision: when SOC 2 Type I scoping commences, when penetration testing is performed, or when the sub-processor list changes — whichever comes first.