SECURITY & AUDIT WHITEPAPER · V1.0

A working framework for auditing AI tools that touch confidential documents.

The four security claims SealedBrief makes about its own design, with step-by-step instructions for verifying each one on the live binary — and ten questions to put to any AI vendor before you trust them with your client files.

AUTHOR	AUDIENCE	DATE	PAGES
SealedBrief engineering team	Lawyers · Therapists · Accountants	May 2026	5

§1 — ABSTRACT

Why this exists.

Every privacy-positioned AI product on the market makes claims about where your data goes and what happens to it. Most of those claims cannot be verified by the buyer. This document is a working framework for changing that.

The framework has two halves. The first is four security claims SealedBrief makes about its own design, with step-by-step instructions for verifying each one on the live binary using tools available on any modern Linux or macOS workstation. The second is a ten-item checklist a buyer can put to any AI vendor — cloud or local — to distinguish honest claims from marketing assertions. SealedBrief is built to satisfy every item; competitors are scored against the same list so the comparison is structural, not promotional.

The regulatory reality

Three professional regimes bind the typical SealedBrief buyer. None of them are abstractions:

Lawyers. Client confidentiality under *ABA Model Rule 1.6* prohibits sharing privileged information with third parties absent informed consent. Cloud LLM providers are third parties. Pasting a deposition transcript into a third-party API routes that material through the vendor's servers, contractors, and logging pipeline — a chain the lawyer cannot audit and rarely fully understands.

Mental health practitioners. PHI shared with a cloud AI provider triggers the Business Associate Agreement (BAA) regime under *HIPAA*'s Privacy and Security Rules. A BAA must be in place; the BAA is enforceable; and most consumer-facing AI tools simply do not offer one. The clinician who pastes a treatment plan into a chatbot to "just rephrase a paragraph" crosses the regime without realising it.

Tax and accounting practitioners. *IRS Pub 4557* and the *Gramm-Leach-Bliley Safeguards Rule* require a written information security program covering client PII (SSNs, financial records, correspondence). Third-party AI tooling triggers vendor-due-diligence obligations the small practice rarely has staff to discharge.

The local-first thesis

Instead of accepting any of the above, run the AI on the practitioner's own machine. The architectural guarantee is "no data leaves the device"; the verification framework in §2 below lets you confirm that guarantee yourself, without trusting the vendor's word for it. The product is local-first by construction, not by promise.

§2 — FOUR CLAIMS, FOUR AUDITS

Verify the design without trusting the vendor's word.

Each claim below is checkable on the live binary in under fifteen minutes. Tools required: `tcpdump` or `nethogs` (network), `ent` (entropy), `gcore` + `gre p` (memory), and the SealedBrief audit harness shipped under `/opt/sealedb rief/audit/`.

01 Zero network traffic from the document-handling process.

NETWORK AUDIT

Contract. Two operating-system processes run when SealedBrief is open: the Presentation Plane talks to our licence server (only) for activation and update checks; the Compute Plane handles your documents and never opens a network connection.

How to audit. Start the app. Identify the Compute Plane PID with `pgrep -af sealedbrief.comp ute`. Run `sudo nethogs -p $PID`. Ask any question. Expect zero non-loopback connections. Anything else is a finding.

02 Encryption at rest is real, not just claimed.

ENTROPY AUDIT

Contract. SealedBrief writes its index, vector embeddings, and metadata to a SQLCipher database under `~/.local/share/sealedbrief/vault/`. The file must look like random ciphertext to a hex dump. We pin the entropy floor at 7.5 bits per byte across every encrypted page.

How to audit. Run `ent ~/.local/share/sealedbrief/vault/core.db`. Expect entropy ≥ 7.5 bits/byte. Optional: open in a hex editor and confirm no ASCII sentences, no `%PDF` headers, no recognisable structure. Below the floor → finding.

03 Master key never leaves the OS keychain.

MEMORY AUDIT

Contract. The AES-256-GCM master key lives in the OS keychain (libsecret on Linux, Keychain on macOS). It enters process memory only for the duration of an active AES context. The key bytes never touch disk outside the keychain itself.

How to audit. While the app is running, dump the Compute Plane memory: `sudo gcore $PID`. Read the key from the keychain: `secret-tool lookup service sealedbrief account master`. Then `grep -ao $KEYBYTES core.$PID`. Expect zero hits outside an active AES context.

04 Every persisted document field is encrypted, not just the obvious ones.

AT-REST AUDIT

Contract. Encrypting "the database" is not enough. Every metadata field that touches user content — file paths, query history, extracted text, OCR output, vector embeddings, timestamps — must run through the field-encryption layer. New ingestion formats fail this gate until they're wired through.

How to audit. Run the per-format coverage tool:

```
python -m sealedbrief.tools.audit_at_rest_encryption \
  --vault ~/.local/share/sealedbrief/vault/ \
  --emit-findings
```

The tool walks every persisted artefact and reports any that fail the entropy check. Empty findings list → claim verified.

§3 — TEN QUESTIONS FOR ANY AI VENDOR

The vendor checklist.

Use the list against any AI tooling you're considering — including SealedBrief. A vendor that can't answer all ten without an NDA, or that needs to ask its sales team, is a vendor whose security posture you can't audit. We've marked SealedBrief's answer next to each item.

01 **Where do my documents go when I ask a question?**
A vendor that says "to our servers" is honest about being a cloud product. A vendor that says "depends on your subscription tier" is a vendor you can't audit. **SealedBrief: never leave your machine.** ✓

02 **Can I monitor the network traffic and verify the previous answer?**
A vendor that can't be put under `tcpdump` fails the audit. **SealedBrief: see §2 claim 01.** ✓

03 **Where does the encryption key live, and who can read it?**
A vendor that says "we manage encryption" controls the key. A vendor that says "your OS keychain" doesn't. **SealedBrief: OS keychain.** ✓

04 **Is the cryptographic primitive named, audited, and standard?**
AES-256-GCM is the modern bar. Vendor-X-proprietary-cipher is not. **SealedBrief: AES-256-GCM, SQLCipher (audited library).** ✓

05 **What metadata gets stored, and where?**
Filenames, query history, embeddings — every field is a potential leak. A vendor that won't enumerate gets a no. **SealedBrief: §2 claim 04 enumerates the layer.** ✓

06 **What logs the vendor keeps, and for how long?**
"We don't keep logs" is uncheckable. "We keep logs for 90 days at provider X" is. **SealedBrief: zero logs of document content; access logs at sealedbrief.com only for licence validation, retained 30 days.** ✓

07 **Does the EULA permit human review of "data we collect"?**
"Authorised personnel may review submissions to improve our models" is a red flag. **SealedBrief: nothing collected, nothing to review.** ✓

08 **What happens when the vendor goes out of business?**
A cloud product becomes inaccessible the day the API turns off. A locally-run product keeps working. **SealedBrief: perpetual licence, offline-verified, no server dependency after activation.** ✓

09 **Can my IT team run a security review without an NDA?**

Vendors who require NDAs to discuss security rule out audits. **SealedBrief: architecture is published; claims are verifiable from the binary without an NDA.** ✓

10 **Opt-in or opt-out for any data collection?**

Opt-in by default with explicit per-channel consent is the privacy-respectful baseline. Opt-out (or no toggle) means your data is the product. **SealedBrief: zero collection by default; opt-in crash-report channel only if you explicitly enable it.** ✓

§4 — CLOSING

If this framework helps, we've succeeded.

Buy SealedBrief if your work involves documents that simply cannot ride a third-party API. Don't buy it if you need agentic workflow execution, cloud-scale inference across many users, or multi-user shared workspaces — the **What SealedBrief is not** section on the website enumerates what we don't do, in five explicit bullets, on purpose: filtering wrong-fit buyers before they buy is part of our design.

A 30-day no-questions refund is in place. If SealedBrief doesn't fit your workflow, the refund process is at `sealedbrief.com/refund`. The licence is added to a revocation list on refund; documents you processed during the window remain yours.

Contact

Technical questions about anything in this whitepaper: **alexandre@sealedbrief.com**.
Refund requests: **refunds@sealedbrief.com**.
Privacy / data-deletion requests: **privacy@sealedbrief.com**.
General support: **support@sealedbrief.com**.

Audit walkthroughs in operating-system terms (Linux today, macOS in V1.0.1, Windows when the SQLCipher dependency stabilises) ship with the product on first install at `/opt/sealedbrief/audit/README.md`.
