

Hidden Tokens against the CNB's AI reading grid

France's National Bar Council (Conseil national des barreaux, CNB) sums up every AI tool it has audited in a ten-field summary form. Here are those ten fields filled in for Hidden Tokens. Every answer points to what lets you check it; where a field does not apply, we say so and say why.

Tardigrade SARL · updated 24 October 2026 · versions described: desktop app 0.8.0, connector 2.22.0, browser extension 0.3.0

IN SHORT


- **Hidden Tokens is not one more legal AI tool.** It is the layer that codes the case file on the lawyer's computer before it reaches whichever AI tool the firm has chosen, then puts the real values back into the result, on the same computer. It doesn't replace those tools: it makes using them defensible.
- **"Where is the data hosted?" It never leaves the workstation.** Hidden Tokens has no server; the publisher receives no file, no key and no log.
- **Two fields are about the LLM and one about RAG** (fields 1, 3 and 7). Hidden Tokens is neither an LLM nor a RAG system: those fields don't apply to it, and we say what the LLM receives instead — codes.
- **Every statement points to a proof:** a check you run yourself on the workstation, a file in the source code (open to audit), or a commitment in the contract.

Why this file

The CNB's [reading grid](#) (1st edition, June 2025, v1.2) ranks four sovereignty criteria; the second is the nationality of the company hosting client data (the fourth applies the same rule to the LLM). It rules out companies whose national law allows extraterritorial requisitions of data, and names US companies under their current legislation as the example. To remove those risks it sees only two routes: full prior anonymisation of the data, or installing the LLM on the firm's own systems. The CNB's guide on ethics and AI (April 2026) widens the first route: it recommends systematically pseudonymising or anonymising data before using any generative AI tool.

Across the grid's fifteen summary forms, by our count, fourteen name a US company among those hosting the LLM (the fifteenth reads "N/A"), and nine declare nothing on pseudonymisation or anonymisation ("N/C" or "N/A"). This is not a criticism of those tools: the large models are, for the most part, hosted by US companies. That is where Hidden Tokens fits. Placed upstream of the AI tool, it ensures that the LLM, wherever it is hosted, only ever receives codes, and that the key to read them stays with the firm.

Where Hidden Tokens sits



Here the CNB form therefore describes two distinct things: the AI tool, its LLM and its servers, which keep their own form; and what that tool receives, which Hidden Tokens reduces to codes. The answers below are about Hidden Tokens. For the AI tool itself, its own form remains the reference.

Two paths lead to the AI tool. **Files** (Word, Excel, PDF, text, scans) are coded into a copy placed in the **Hidden-Tokens/** subfolder; that copy is what the AI gets, through Claude Desktop or by uploading it to any tool. **Typing online** is coded by the browser extension on ChatGPT, claude.ai, Gemini and Copilot, and only there for now.

The summary form

1 LLM HOSTING LOCATION

Not applicable: Hidden Tokens is not an LLM and hosts none. The chosen tool's LLM only receives codes.

2 DATA HOSTING LOCATION

The user's computer; the shared registry, encrypted, on the firm's own file share. No Hidden Tokens server.

3 COMPANY HOSTING THE LLM (NATIONALITY)

Not applicable: it is the AI provider the firm chooses (Anthropic, OpenAI, Google, Microsoft: all US).

4 COMPANY HOSTING THE DATA (NATIONALITY)

No third party: the firm itself. The publisher, Tardigrade SARL, a Moroccan company, receives no data.

5 REUSE OF CLIENT DATA BY THE LLM

What the LLM could reuse is reduced to codes. Local proof: the "What Claude saw" log.

6 DATA ENCRYPTION

Shared registry in AES-256-GCM, workstation keys in the OS keychain, key held by the firm alone. Single-workstation key: account permissions, no encryption of its own.

7 SECURE RAG

Not applicable: Hidden Tokens builds no index. A tool that indexes the coded copy only indexes codes.

8 PSEUDONYMISATION / ANONYMISATION

Reversible pseudonymisation, on the computer, before the AI: it is what the product does. Benchmark: 201 of 201 sensitive values coded.

9 CLIENT DATA ACCESSIBLE TO STAFF

No access possible for the publisher's staff: nothing reaches it. Support works on codes.

10 ADDITIONAL MEASURES

Log of what the AI received, signed proof pages, a seal the client can verify, coverage report, signed policy, notarised app, no telemetry.

Dashed boxes: field not applicable to Hidden Tokens, explained below.

Field by field, with the proof

ON YOUR COMPUTER a check that you, your DPO or your IT team run yourselves, without taking our word for it;

IN THE CODE a file in the source code or technical documentation, available for inspection during an audit (framework agreement, article 11);

IN THE CONTRACT a written commitment by the publisher in the Organisation framework agreement (French).

1 LLM hosting location

Not applicable. Hidden Tokens is not a language model and hosts none, on the computer or anywhere else.

The LLM is that of the AI tool the firm chooses; it stays hosted wherever that tool's form says. Hidden Tokens doesn't move that LLM: it changes what the LLM receives. The text that reaches it carries codes (**NOM-001** , **SOCIETE-003**) in place of names, companies and identifiers.

PROOF

ON YOUR COMPUTER Turn off the internet: coding and decoding still work. Hidden Tokens calls no model.

IN THE CODE `test/reseau.test.js` (connector): the whole flow — spreadsheet, PDF, OCR of a scan, batch, log, dashboard — runs with networking forbidden by Node's permission model; any connection would fail.

LIMIT Hidden Tokens does not remove the need to check where the chosen tool's LLM is hosted: it reduces what that LLM sees.

2 Data hosting location

The data never leaves the workstation. Files in clear, the mapping key, decoded files and the log stay on the user's disk; a firm's shared registry lives on its own file share (NAS, file server, SharePoint), encrypted.

Hidden Tokens has no server: no account, no online storage, no licence call. The licence is a signed token, verified on the computer.

PROOF

ON YOUR COMPUTER With an application firewall (LuLu, Little Snitch), the connector attempts no outgoing connection.

ON YOUR COMPUTER The desktop app's only outgoing connection is the update check on GitHub (in version 0.8.0: at launch, then every six hours). It sends the IP address and the installed version; no file, code or key.

IN THE CODE `extension/manifest.json` : the extension's only network permission is `http://127.0.0.1`, the app on the same computer. Its scripts run in the ChatGPT, claude.ai, Gemini and Copilot pages to code the message before it is sent ([extension privacy policy](#)).

IN THE CODE `server/licence.js` : licence verified by Ed25519 signature, offline. `server/registre.js` : the shared registry is a folder on the firm's share; no server.

IN THE CONTRACT Article 8: the software runs with no connection to the publisher, no account, no telemetry. Article 9: it sends nothing to the publisher.

LIMIT A working folder synced to a cloud (iCloud, Google Drive, Dropbox, OneDrive) takes the files and the key off the computer. Hidden Tokens detects and reports it (containment audit, `verifier_confinement` tool); it does not prevent it.

3 Company hosting the LLM (nationality)

Not applicable to Hidden Tokens: it is the company of the AI provider the firm chooses.

For the most widely used assistants — Claude (Anthropic), ChatGPT (OpenAI), Gemini (Google), Copilot (Microsoft) — these are US companies, precisely the ones the grid's second and fourth sovereignty criteria target. Hidden Tokens doesn't change their nationality. It ensures that none of them receives a client's name, and none holds the key that would let them find it.

PROOF

ON YOUR COMPUTER The "What Claude saw" log (field 5) shows the exact text given to the AI and re-checks it against the key.

IN THE CODE `server/sortie.js` : on the final text of every tool response, any value known to the key turns back into its code, and any sensitive pattern (email, phone, IBAN, long identifier) is masked.

LIMIT The status, under the GDPR and transfer rules, of pseudonymised data given to a recipient without the key is for the firm and its counsel to assess (EDPB guidelines 01/2025; CJEU, 4 September 2025, C-413/23 P).

4 Company hosting the data (nationality)

No third party: the data stays with the firm. The publisher, Tardigrade SARL (Casablanca, Morocco, RC 620437), receives, stores and processes none of it.

The publisher therefore has no file, key or log to hand over to anyone. For the same reason it is not the firm's processor under article 28 of the GDPR for data handled by the software: the firm remains the sole controller.

The software itself is distributed through GitHub (Microsoft, United States) and this website is hosted by Netlify (United States). Those channels carry the signed software and public pages, never case-file data.

PROOF

IN THE CONTRACT Article 9: the publisher receives, stores and processes no client data — no registry, no mapping key, no log — and does not act as a processor (GDPR art. 28; Moroccan law 09-08).

ON YOUR COMPUTER The network checks in field 2.

LIMIT The publisher does process business-relationship data (contacts, invoicing), described in its [privacy policy](#).

5 Reuse of client data by the LLM

Hidden Tokens cannot bind the LLM provider on what it does with what it receives. It ensures that what it receives, and could therefore reuse, is reduced to codes — and proves it on the computer.

The CNB's guide observes that a provider's promise not to reuse data is generally neither provable nor verifiable. Hidden Tokens doesn't ask you to trust a promise: the local log shows what left.

PROOF

ON YOUR COMPUTER Ask Claude "show me what you saw" (`journal_ia` tool): the local page `CE-QUE-CLAUDE-A-VU.html` lists every text Hidden Tokens gave the AI and the number of real values found in it. Expected result: 0.

IN THE CODE `server/journal.js` : every response is logged after the output lock, i.e. exactly what the model receives. Each time the log is opened, all of it is re-checked against the current key: a value that left before the key knew it is found after the fact. A leak is shown by its code, never its value.

IN THE CODE Messages typed into ChatGPT, claude.ai, Gemini or Copilot: the coded text sent is logged the same way; of the clear text, only its length is recorded (`SPEC-VERROU-UNIVERSEL.md`).

LIMIT The LLM receives the codes and the structure of the file — dates, headings, reasoning, non-sensitive amounts — and may use them as its terms allow.

6 Data encryption

Hidden Tokens replaces rather than encrypts: what goes to the AI doesn't contain the values. What it shares and what it keeps is protected as follows, and the key is held and managed by the firm alone.

PROOF

IN THE CODE **Shared registry**, on the firm's share: every entry encrypted with AES-256-GCM; secrets sealed for each certified workstation (X25519, HKDF-SHA256, AES-256-GCM); the registry signed with Ed25519 by the administrator. The share doesn't need to be trusted (`server/registre.js` , `server/enveloppe.js`).

IN THE CODE **Each workstation's private key**: the login keychain on macOS, DPAPI on Windows; never a silent fallback to a plain file (`sentinelles/lib/trousseau.js`).

IN THE CONTRACT **Key custody**: by the firm. The escrow opens only with a quorum of guardians appointed by the firm; a recovery sheet can be printed. The publisher holds no key and can restore nothing (article 8).

IN THE CODE **Ed25519 signatures**: the firm's policy, the licence, the seal, the proof pages.

ON YOUR COMPUTER **In transit**: towards the AI, HTTPS/TLS encryption is the AI provider's. Between the extension and the app, traffic goes through `127.0.0.1` and never leaves the computer.

LIMIT A workstation's mapping key (`cle-de-session.json`) is protected by account permissions (readable by that user only), not encrypted by Hidden Tokens in this version (`server/index.js`). Turn on disk encryption (FileVault on macOS, BitLocker on Windows): that is what protects it at rest.

7 Secure RAG

Not applicable. Hidden Tokens builds no index, no vector store and no document search engine, on the computer or anywhere else.

If the firm's AI tool indexes documents, it can be given the coded copy Hidden Tokens produces: its index will then hold only codes, and a Word document it produces from that copy is decoded on the computer (`decoder_word` tool). The security of the RAG itself remains the tool's.

PROOF

IN THE CODE The connector's sixteen tools code, decode, check or report; none indexes. The connector writes only the key, the log and the coded copies, in `Hidden-Tokens/` .

8 Pseudonymisation / anonymisation

Reversible pseudonymisation, done on the computer, before the AI. It is what Hidden Tokens does.

Names, companies, addresses, identifiers (ID card, passport, SIREN and SIRET company numbers, IBAN, court docket, prosecution and appeal numbers...), references and, in a spreadsheet, amounts that reveal a person's finances (pay, balances, fees, adjustable column by column) become codes. The key is cumulative: the same person gets the same code in every document of the case file, and with the shared registry across the whole firm, so the AI can reason over the whole file. Pseudonymisation is deliberate: only it lets the firm hand the client work that carries the real names. The CNB's guide expressly accepts it, alongside anonymisation.

PROOF

ON YOUR COMPUTER Coding happens in two steps: the plan is shown (what will be coded, what will stay readable), the user approves it, then the copy is written. Open it: no client name.

IN THE CODE Three stacked locks. The lock (`server/verrou.js`) leaves readable only what is proven harmless — words from a 336,000-form French dictionary, non-identifying numbers, an allow-list (courts, codes, document titles) — and codes everything else. The preview check re-runs every cell shown to the AI. The output lock (`server/sortie.js`) re-runs every response.

IN THE CODE Re-identification test (`server/reidentification.js`): on the coded text, it counts indirect clues found together in one passage (place, age, rare job title, date of birth); dates of birth are coded by default; the `generaliser` option codes cities and turns an age into a range.

IN THE CODE Benchmark (connector 2.10.0, 21 October 2026): 201 of 201 sensitive values coded across 15 fictitious documents (law, finance, payments), 0 leaks (`bench/resultats.json` , `npm run bench`). Leak hunt across every tool: `test/fuite.test.js` .

LIMIT Pseudonymising is not anonymising: for the firm, which holds the key, the data remains personal data (GDPR art. 4(5)). A word that is first a city, a day or a role, used alone as a surname ("Paris"), may stay readable; indirect clues are flagged, not all rewritten. Hence the review of the plan.

9 Client data accessible to staff

No access possible for the publisher's staff: no client data ever reaches it.

Support works on codes: to report a problem, the user quotes `NOM-012`, never the value, and the publisher never asks for client data. The software has no remote access and no publisher console on the workstations.

Within the firm, adding or removing a workstation, opening the escrow and any restore are logged and signed in the registry; the coverage report gives counts only.

PROOF

IN THE CONTRACT Article 7: the client quotes codes, never values; the publisher will never ask for client data. Article 9: no data at the publisher.

ON YOUR COMPUTER The network checks in field 2: nothing goes to the publisher.

IN THE CODE `SPEC-REGISTRE-ORGANISATION.md`, invariant 5: adding or removing a workstation, opening the escrow and restoring are logged in `registre/journal.jsonl`, signed.

LIMIT In this version, every certified workstation on the shared registry can decode all of the firm's codes: partitioning by team or by case file is not yet available.

10 Additional measures

ON YOUR COMPUTER **Signed proof pages.** The log and the coverage report carry a token signed (Ed25519) over their counts and the page's SHA-256 fingerprint: a figure changed by hand breaks the signature. Verified offline in the app ("Verify a proof...").

ON YOUR COMPUTER **The seal for the client** (`sceau_dossier` tool): signed proof that their file went through Hidden Tokens without a leak, which the connector refuses to issue if it finds one. The client verifies it on hiddentokens.com/verify-seal, offline: the seal travels in the part of the address the browser never sends to the server.

IN THE CODE **The coverage report** (`rapport_couverture`, `server/couverture.js`): workstation by workstation, what was protected before the AI and which machines no longer go through Hidden Tokens. Counts only, never a value, a code or a file name.

ON YOUR COMPUTER **The containment audit** (`verifier_confinement`) names the other doors into the folder: another extension or MCP server with access to it, Cowork mode, browser or screen control, a cloud-synced folder.

IN THE CODE **The firm's policy is signed** (Ed25519) and re-checked at every coding: missing, altered or rolled back, coding stops (fail closed). Decoding is never blocked (`server/politique.js`).

IN THE CODE **The extension holds the message back** once it has been paired: app not responding, text not replaced in the input box, blocking rule — nothing leaves (`extension/README.md`). An organisation can enforce this from installation through Chrome policy.

ON YOUR COMPUTER **App signed and notarised by Apple** (Developer ID, team 2CTJY8T842): `spctl -a -vv "/Applications/Hidden Tokens.app"` answers "source=Notarized Developer ID". Every published version carries its SHA-512 checksums ([releases page](#)).

IN THE CONTRACT **No telemetry**, no usage statistics (article 8). **The code is open to audit** by the client or an independent firm it appoints, together with the proof verifier and the benchmark (article 11).

IN THE CODE **Scans are read on the computer**: text recognition is built in (`server/ocr.js`), with no online service.

IN THE CODE **Documents provided**: pre-filled data protection impact assessment template, AI policy template, security sheet, mapping to the GDPR, Moroccan law 09-08 and the CNB guide (the connector's `confiance/` folder, in French).

LIMIT Hidden Tokens has, to date, no certification, no label and no published independent security audit. The Windows installer is not yet code-signed: Windows shows "unknown publisher" at first install. Before it has ever been paired, the extension lets messages through, flagged by a red banner.

Check it yourself in ten minutes

1. Install the app (14-day trial, no account) and check its signature: `spctl -a -vv "/Applications/Hidden Tokens.app"` .
2. Turn off the internet. Code a fictitious case file, have Claude work on it, decode the Word file produced: it all works.
3. Turn the internet back on behind an application firewall (LuLu is free): the connector attempts no connection; the app contacts only GitHub, for updates.
4. Open the coded copy in `Hidden-Tokens/` : no client name.
5. Ask Claude "show me what you saw": the local page shows 0 real values.
6. Ask "check containment": you get the list of other doors open into the folder, and the step that closes each one.
7. Issue a seal and verify it on hiddentokens.com/verify-seal with the developer tools' network tab open: the seal appears in no request.

What Hidden Tokens does not cover

- Text typed or pasted straight into a chat without the extension, a file attached by hand: they go through no lock. The firm's AI policy makes Hidden Tokens the only authorised route; the coverage report lets you check it is followed.
- Typing online anywhere other than ChatGPT, claude.ai, Gemini and Copilot, including a legal AI tool's own interface: you upload the coded copy there, but what you type is not coded.
- Indirect clues (job title, age, context): they are flagged, not all rewritten.
- Tasks where pseudonymisation doesn't apply, such as a live summary of a client meeting, which the CNB's guide cites.
- The firm's own obligations: impact assessment, choosing and contracting the AI provider, informing the client. Hidden Tokens supports compliance; it doesn't certify it.

Sources

- CNB, Grille de lecture — synthèse des consultations effectuées par le groupe de travail sur l'intelligence artificielle – <https://cnb.avocat.fr/medias/cnb-grille-lecture-intelligence-aartificielle-2025-68f775b81ba5b2.65545272.pdf>, 1st edition, June 2025, v1.2 (French).
- CNB, commission des règles et usages, *La déontologie et l'intelligence artificielle*, 1st edition, April 2026 (French).
- GDPR, articles 4(5), 25, 28 and 32; EDPB guidelines 01/2025 on pseudonymisation; CJEU, 4 September 2025, EDPS v SRB, C-413/23 P.
- Hidden Tokens, Hidden Tokens Organisation framework agreement – [/fr/contrat-cadre/](#) (French), articles 7, 8, 9 and 11.
- Hidden Tokens, extension privacy policy – <https://hiddentokens.com/extension/privacy/>; published versions and checksums – <https://github.com/ai4xedu/hidden-tokens-releases/releases>; verify a seal – [/verify-seal/](#).

Hidden Tokens supports your compliance; it doesn't certify it. This file describes the software; it is not legal advice. © 2026 Tardigrade SARL. Claude is a trademark of Anthropic; ChatGPT, Gemini and Copilot are trademarks of their respective owners. Hidden Tokens is not affiliated with any of them, nor with the Conseil national des barreaux.