NOCIP

No Cloud Inference on Personal data

Audit Report

LexITanon Desktop (Linux) · 0.5.0 (Linux x86-64)

Label awarded on 2026-11-10: NOCIP Gold · index 10/10 · No. G-2026-0001

1 Details

Report number	NOCIP-PB-2026-0001
Label ID	NOCIP-G-2026-0001
Vendor	be next GmbH, Kirchheimer Str. 27, 70619 Stuttgart
Product	LexITanon Desktop (Linux)
Version	0.5.0 (Linux x86-64)
Tier applied for	Gold
Criteria catalogue	1.0 (freigegeben am 11.11.2026 / released 2026-11-11)
Test corpus	NOCIP test corpus v0.1: 15 documents (German, English), 168 canary values
Audit period	2026-10-29 to 2026-11-11
Auditor	be next GmbH (Prüfstelle NOCIP)

Scope: The audit covers the Desktop version for Linux (one workstation, access from the same computer only), first in version 0.3.0 (2026-10-29), and again for the award in the stable version 0.5.0 (2026-11-10). The program file is the same for Desktop and Server; the label covers the Desktop version. The Windows, macOS and Server versions are not covered by this report: a capture on Windows was not possible, and the Server version does not meet G-02 (finding 6). The label covers later versions once NOCIP has checked them before release (C-05).

be next GmbH is both the vendor of this product and the auditor. Under the NOCIP rules, the audit follows the published audit checklist in full, and this report is published in full.

2 Result

LexITanon's AI runs entirely on the customer's computer: its own recognition model, shipped encrypted inside the program and checked when loaded. Processing the test corpus caused not a single connection attempt; the program works fully without network. The only connections are the licence check, updates on click and “report a wrong recognition”, which sends only with consent and only invented data. The licence agreement (version 2 of 2026-11-09) states that data processed with LexITanon is not transferred to be next.

All Gold criteria are met. The label NOCIP-G-2026-0001 is awarded for LexITanon Desktop (Linux) 0.5.0: Gold, NOCIP index 10 of 10 (groups I-1 and I-5 as worded in the catalogue of 2026-11-11), valid until 2027-11-09. The disclosure on lexitanon.com and in the manual (C-01) follows with the entry in the register.

State of the criteria	
passed	10

State of the criteria	
not applicable	1

3 Technical test

Test	Result	Finding
A · Checksum	passed	Program file lexitanon-0.5.0-linux-x64: SHA-256 9570b9ed…f2f32 = signed release manifest (stable since 2026-11-08) = published checksum. Likewise on 2026-10-29 for 0.3.0 (ffc13f42…dcc8de6).
B · Self-test without network	passed	In a network namespace without network: recognition rules version 4, fonts, model 1.3.3 (embedded, 37 labels, ONNX Runtime 1.31.0) loaded and checked - “self-test passed”.
C · Processing under capture, without network	passed	The program ran in a namespace without any connection outside; every connection attempt was recorded with strace. 15 documents of the test corpus through /redact-text and one job through /jobs: all processed. Connection attempts: 0 (not even name resolution). “Report a wrong recognition” with and without consent: HTTP 503 (“only possible with the online licence”), no connection. Licence: the website's demo licence (at most 20 findings per document) - it limits recognition, not network behaviour. Same result on 2026-10-29 with 0.3.0.
E · Canary values in the program's own files	passed	Program log and session log after processing: none of the 168 canary values present.
F · “Report”: invented data in line 2	passed	The replacement in web/feedback.js (in the browser, without a mapping table; the file is byte for byte the one embedded in 0.5.0) with the 15 documents: all 168 canary values replaced, 0 let through. In 0.3.0 it was 167 of 168 (one e-mail address with an umlaut); fixed since 0.3.1.
G · Screenshots	passed	The interface in the same isolated namespace (section Screenshots): result of a test document, “report a wrong recognition” with lines 1 and 2 and the consent checkbox (“Send” active only once ticked), the answer on sending. The demo licence starts the program without its interface; the project's interface files were therefore served through a small relay to the program's API. Five of the six files, feedback.js among them, are byte for byte the ones embedded in 0.5.0 (SHA-256 per its start-up check); app.js was changed after the release.

4 Results per criterion

C-01 Disclosure

Requirement	The product and the vendor's website state the NOCIP tier and say where AI processing takes place. Users can find this statement before first use (e.g. product page, settings, documentation).
Result	**passed**
Finding	Website and manual say “entirely local, no cloud” and describe the connections precisely (lexitanon.com: home page, features, AI & professional law; manual ch. 2.1 to 2.3), since 0.4.0 including “report a wrong recognition” (only on request, only invented data). With the award, lexitanon.com and the manual state “NOCIP Gold · No. G-2026-0001” with a link to the verification page in the register (finding 1).

C-02 Model inventory

Requirement	The documentation lists every model used for the product's AI functions: name, version, licence, origin and where it runs (customer, vendor or third-party AI, with country).
Result	**passed**
Finding	Manual ch. 2.2, table “Das KI-Modell im Überblick”: own recognition model (token classification, BERT), version 1.3.3 with recognition rules 4 (from 0.5.0), base model: German BERT base model (deepset GmbH, MIT licence), model licence proprietary (be next GmbH), run locally on the CPU with ONNX Runtime 1.31, built into the program file and checked on start. Consistent with test B.

C-03 No training on customer data

Requirement	Customer content is not used by the vendor or any third party to train or fine-tune models unless the customer explicitly opts in. Opt-in is off by default. Fine-tuning that takes place only on the customer's own systems is allowed.
Result	**passed**
Finding	Licence agreement version 2 of 2026-11-09 (lexitanon.com/agb/): “personal data processed with LexITanon is not transferred to be next, and be next has no access to such data” (cl. 13.2); be next receives only the licence-check data and what the customer provides voluntarily (cl. 12.1); the usage statistics contain no content (cl. 6.3). Manual ch. 2.2 and website: the model does not learn from the customer's documents. The only exception are examples from “report”: only with explicit consent for every report (checkbox, empty every time; text “… used to improve the recognition of LexITanon”); without consent, program and licence server refuse (web/feedback.js, src/api.cpp, lexlic.php), and they contain only invented data (test F).

C-04 No content in telemetry and logs

Requirement	Telemetry, usage statistics, crash reports and remote logs contain no document content, prompts, AI output or personal data taken from processed content.
Result	**passed**

Finding	The licence check sends the licence key, installation ID, hashes of hardware features, platform and version, but no content (src/licence_online.cpp, src/machine_id.cpp). No crash reports; ONNX Runtime telemetry is switched off (src/model_ort.cpp; tools/build_ort_static.sh). The session log on the computer contains job number, size, number of findings and time, no text (src/jobs.cpp). The usage counts (since 0.3.0) contain no content either (src/usage.cpp; licence agreement cl. 6.3). Tested (tests C, E): nothing transferred during processing, no canary values in logs.

C-05 Change notification

Requirement	The vendor reports changes to the AI architecture (new model or model provider, new processing location, new external connection, changed router or pseudonymisation) to NOCIP within 30 days. A change that leads to a lower tier is reported before it goes live.
Result	**passed**
Finding	The auditor's own product: instead of the label agreement, be next commits to naming changes to the AI architecture in the release notes and to having NOCIP check them before release. From 0.3.0 to 0.5.0, new models (1.3.1, 1.3.3) and the sending of “report” were added; both are re-checked in this report.

C-06 Connection list

Requirement	The vendor lists every outgoing network connection the product makes (destination, purpose, data sent). Processing the test documents creates no connections outside this list.
Result	**passed**
Finding	Connections according to the source code: https://lexitanon.com/key/api/v2/ (activate, check, release, download, feedback), as a fallback https://key.lexitanon.com/api/v2/; Desktop once a day; a proxy from https_proxy is used (src/licence_online.cpp). Updates only when the user clicks, “report” only at the user's request. Described in manual ch. 2.3 and on “AI & professional law”. Tested (test C): not a single connection during processing.

C-07 Data processing agreement

Requirement	If the vendor processes customer data (e.g. SaaS, remote support), it offers a data processing agreement under GDPR Art. 28 that names all sub-processors with access to AI input or output. Does not apply if the vendor receives no customer data.
Result	**not applicable**
Finding	be next receives no document content (licence agreement cl. 12.1, 13.2) and is not a processor (cl. 13.2). Reports from “report” contain only invented data. be next processes the licence-check data for the contract itself; the privacy policy does not describe it yet (finding 4).

C-08 User-triggered reports

Requirement	If the product sends examples to the vendor at the user's request (e.g. to improve recognition), then only like this: (1) the user triggers each report, never automatically and never as a default. (2) The product first replaces personal data with invented data of the same format, without a mapping table. (3) A preview shows exactly what is sent, and the user can change it; words the user keeps in the original are confirmed once more before sending. (4) The vendor uses reports only to improve the product and does not pass them on. Does not apply if the product sends no reports.
Result	**passed**
Finding	Since 0.4.0, “report a wrong recognition” sends. (1) Only the user triggers a report, with a tick for every report, never preset. (2) In the browser, the program replaces names, places, dates, numbers, IBANs, e-mail and web addresses with invented data of the same format; the real line 1 stays on the computer, there is no mapping table. Test F: 0 of 168 canary values let through (thresholds: all groups 100 %). (3) Line 2 is the preview and exactly the text sent (web/feedback.js); the user can keep or replace words, kept words are confirmed before sending. The program forwards only kind, line 2, expected kind, versions and consent (src/api.cpp). (4) Purpose “to improve the recognition of LexITanon” in the consent text; that reports are not passed on is stated with the disclosure on “AI & professional law” (finding 1). A sent report was not captured (the demo licence does not send, the online licence needs the licence server); source code and preview were checked.

G-01 Local inference

Requirement	All inference for AI functions runs on the customer's devices or servers (own hardware, or a cloud environment the customer controls itself). No AI request goes to the vendor or to third parties.
Result	**passed**
Finding	Recognition runs inside the program (ONNX Runtime, CPU, no network access during inference; STRATEGY.md, src/model_ort.cpp). There is no AI request to be next or third parties. Tested (test C): test corpus fully processed without network, no connection attempts.

G-02 Offline operation

Requirement	All AI functions work without an internet connection. Licence checks and updates are allowed if they carry no content and processing continues offline for at least 14 days.
Result	**passed**
Finding	Recognition needs no internet. If the licence server cannot be reached, the Desktop version keeps working for 14 days (licence-web/lexlic/lexlic.php; FAQ on lexitanon.com) - exactly the threshold. Offline updates are possible with --update-from (manual ch. 12). Tested (tests B, C): self-test and processing without network; the grace period was checked in the licence server's configuration (allowed by audit checklist G-02, step 3).

G-03 Verifiable models

Requirement	Models and model updates are delivered as files with a published checksum or signature.
Result	**passed**

Finding	The model is embedded in the program, encrypted with AES-256-GCM and checked on loading through the GCM tag and a SHA-256 checksum; on start the program checks its manifest and a test sentence (src/resources.cpp, src/startup.cpp). New models come only with a new program version; its manifest is signed with Ed25519, the SHA-256 checksums are listed under “My licences”. Tested (tests A, B): checksum recomputed, self-test passed.

The criteria of the Silver and Bronze tiers do not apply: Gold was applied for.

5 NOCIP index

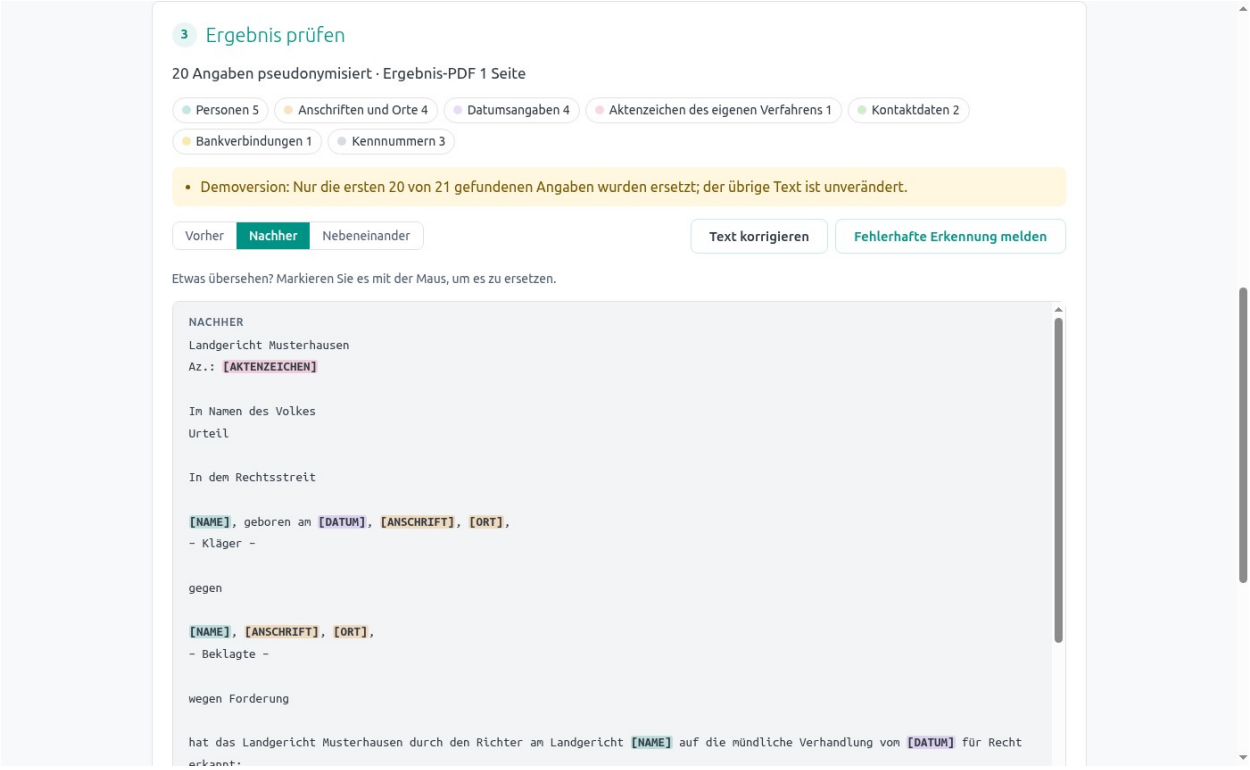
Group	Points	Reason
I-1 Transparency	2	Website (AI & professional law, data-flow diagram) and manual ch. 2.2 state the processing location: everything on the customer's computer. The program has one AI function (recognition) and permanently shows “Läuft auf diesem Rechner” (runs on this computer), when loading “Ihr Dokument bleibt auf diesem Rechner”; the report window says what goes to lexitanon.com.
I-2 Customer control	2	Gold: all AI functions run at the customer's site.
I-3 Data minimisation	2	No content is kept outside the customer; the session log contains no text. Reports contain only invented data.
I-4 Integrity	2	The model is checked before loading; documents are processed in memory only, there are no stores holding content.
I-5 Verifiability	2	Manual ch. 2.3 lists the destinations (lexitanon.com and key.lexitanon.com, HTTPS, port 443), the purpose (licence check with usage counts, downloading a new version, “report”) and when they take place (Desktop once a day, Server hourly; update and report only on click). Checked with test C. Recommendation: add the estimated data volume per day.
Total	**10 / 10**	

6 Findings and actions

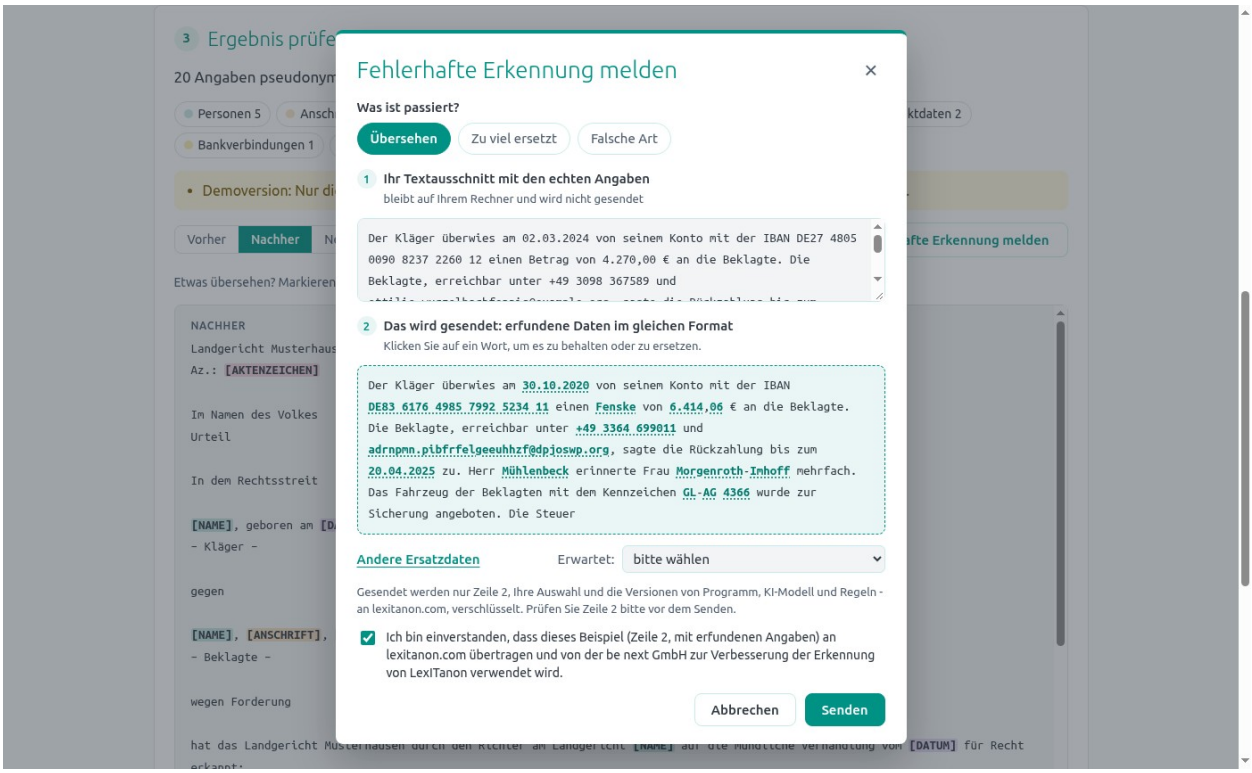
1. Disclosure (C-01, C-08): lexitanon.com and the manual state “NOCIP Gold · No. G-2026-0001” with a link to https://nocip.com/register/NOCIP-G-2026-0001/ and a note that it covers the Desktop version (Linux). On “AI & professional law”, be next adds for “report”: reports are used only to improve LexITanon and are not passed on.
2. Training (C-03): the licence agreement excludes transferring processed data to be next and so covers training. Recommendation for the next version: state expressly that customer documents are not used for training and reports only with consent.
3. “Report”: the replacement rules also replace some nouns with names (e.g. “einen Betrag” → “einen Fenske”, screenshot 2). Harmless for privacy, but it makes the reports harder to evaluate. The thank-you message after sending still names beA.expert in feedback.js.
4. Privacy policy of lexitanon.com: does not yet describe the licence check, the usage counts and “report”; the licence agreement refers to “AI & professional law” for them, where they are described.
5. Windows and macOS: not covered by this label (no capture possible on Windows; macOS not audited). For a later audit: check the connections of WebView2 or WebKit in the capture; on Windows, the bundled onnxruntime.dll is not checked before loading.

6. Server version: keeps working without the licence server for only 24 hours - G-02 (14 days) is not met. Possible: Silver (S-a, the AI runs at the customer's site), or Gold with a longer grace period or an offline licence.
7. Index rescored on 2026-11-11: catalogue 1.0 rewords groups I-1 and I-5 (version of 2026-11-11, released the same day). I-1 rates the processing location in documentation and product (a permanent notice is enough if all AI functions run in the same place); I-5 rates documented connection destinations, purpose and frequency, which the customer can check with their own tools. A connection log in the product and permanent test access are now only recommendations. Index before 7/10 (I-1: 1, I-5: 0), now 10/10.
8. Manual: ch. 2.3 and the footer still name beA.expert instead of be next or LexITanon.
9. History: first audit on 2026-10-29 with 0.3.0 (technically passed; open were the model inventory, the contract, e-mail addresses with umlauts in “report”). Model inventory in the manual since 2026-10-29, umlaut error fixed in 0.3.1, licence agreement version 2 online since 2026-11-10, re-check of the stable version 0.5.0 on 2026-11-10.

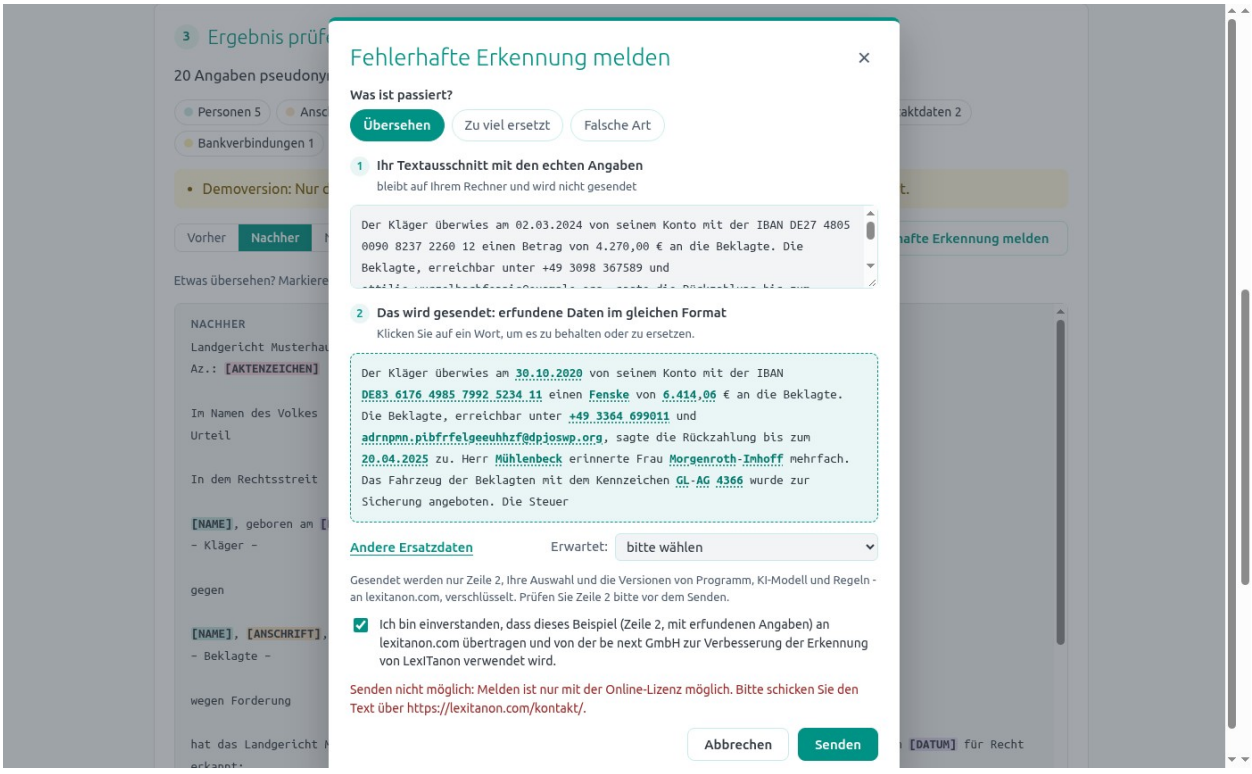
7 Screenshots



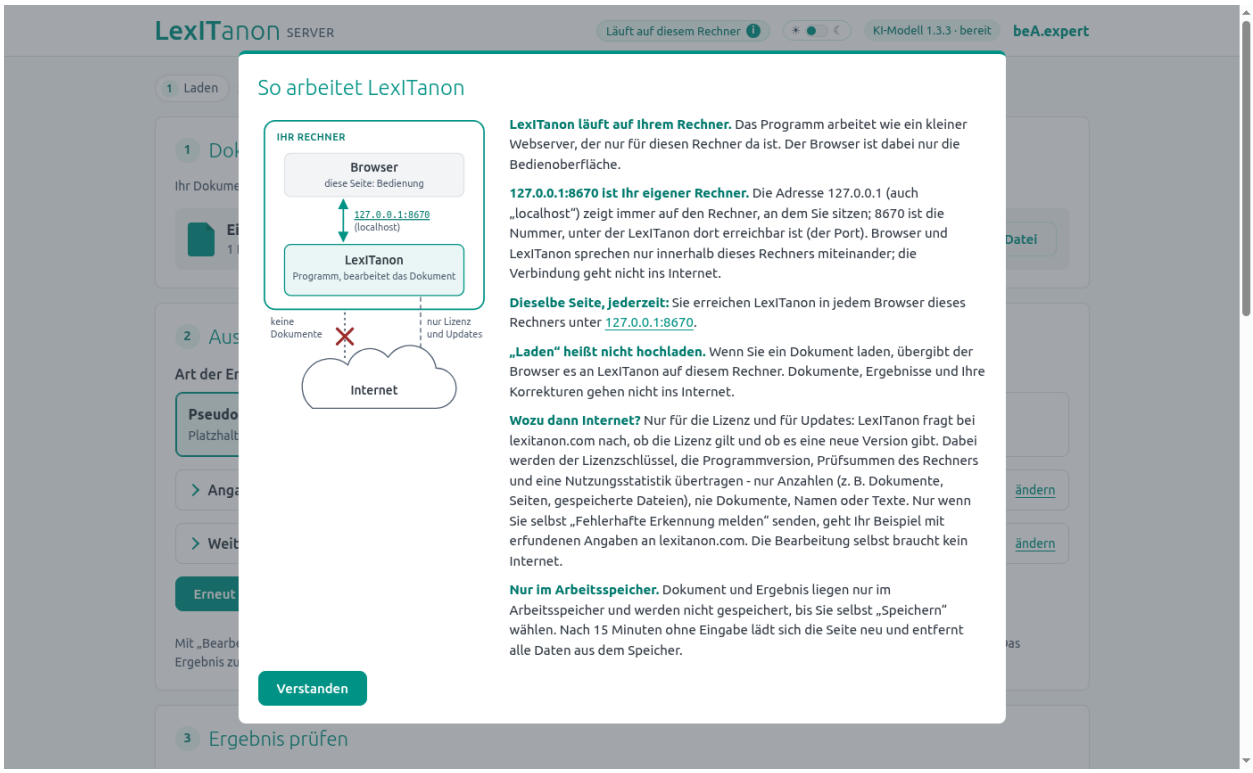
1 · Result of a test document (pseudonymised), processed without network (0.5.0).



2 · “Report a wrong recognition”: line 1 with the real canary values stays on the computer; line 2 with invented data of the same format would be sent, only after the consent is ticked.



3 · After “Send” with the demo licence: nothing sent (“only possible with the online licence”).



4 · "How LexITanon works": the explanation in the program itself.

8 Sign-off

The audit is completed. The label NOCIP-G-2026-0001 (Gold, NOCIP index 10/10) is awarded and entered in the register: https://nocip.com/register/NOCIP-G-2026-0001/. The register shows the current status.

Place, date	Stuttgart, 2026-11-11
Auditor	be next GmbH
Signature	