NOCIP

No Cloud Inference on Personal data

Criteria Catalogue

Requirements and audit procedure for the NOCIP label

Version 1.0 · 2026-10-29 · Draft – not yet approved







Tier	Name	Meaning
Gold	Local AI	The AI runs at the customer's site. Real content never leaves the customer.
Silver	Hybrid AI	Personal data is seen only by AI run by the customer or by the vendor itself.
Bronze	Pseudonymised	A third-party AI receives personal data only in pseudonymised form.

1 Introduction

This catalogue sets out what a product must meet for the NOCIP label and how be next GmbH audits it. It applies to software installed at the customer's site and to software as a service (SaaS).

Every criterion has a fixed ID: C for criteria of all tiers, G for Gold, S for Silver, B for Bronze. A label always names the catalogue version it was audited against.

The tier answers one question: who runs the AI that sees personal data? In addition, the NOCIP index rates the implementation within the tier.

2 Definitions

Term	Meaning
AI function	Any product function whose result is produced by a machine-learning model, e.g. language models, name recognition, classification, speech recognition.
Inference	Running a model on an input to produce a result.
Customer	The organisation that uses the product (usually the controller under the GDPR).
Vendor	The company that offers the product and applies for the NOCIP label.
Vendor-operated infrastructure	Servers the vendor operates itself or rents for its exclusive use, on which the vendor installs and runs the models itself (e.g. open-weights models on dedicated servers or its own GPU instances).
Third-party AI	An AI model that neither the customer nor the vendor runs itself and that is reached through a service. This includes a cloud provider's model services in the vendor's own account, even when hosted in the EU.
Personal data	Any information relating to an identified or identifiable natural person (GDPR Art. 4(1)).
Pseudonymisation	Replacing personal data with placeholders so that the recipient cannot attribute them to a person without the separately kept mapping table (GDPR Art. 4(5)). Pseudonymised data is not anonymous.
Router	A processing step that decides for each request whether it contains personal data and, therefore, where it may go.
Canary value	An invented piece of personal data (name, IBAN, address …) that NOCIP plants in test documents in order to find it again in outgoing traffic.
NOCIP test corpus	German and English test documents and test requests containing canary values. It is not published and is renewed every year.

3 Tiers

Gold · Local AI

All AI functions run on the customer's devices or servers. Neither the vendor nor any third party receives content – except reports the user triggers and that contain only invented data (C-08). The AI functions work without an internet connection.

Silver · Hybrid AI

AI that processes personal data runs at the customer's site or on infrastructure the vendor operates itself. A third-party AI receives, at most, requests without personal data. Silver has two variants; a product may meet both.

- **S-a · Vendor-operated AI:** Software as a service (SaaS) whose models the vendor runs itself.
- **S-b · Local AI with a router:** A router sends only requests without personal data to a third-party AI; all others stay local or with the vendor.

Bronze · Pseudonymised

A third-party AI is used. Before every transfer, a pseudonymisation step replaces personal data with placeholders. The mapping table never leaves the customer or the vendor.

4 Criteria

4.1 All tiers

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).
Evidence	Addresses (URLs) and screenshots of the statements.
Audit	1. Open the statements on the website and in the product. 2. Compare them with the architecture found in the audit.
Passed if	Statements present, accurate and matching the tier.

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).
Evidence	Model inventory (table in the self-assessment).
Audit	1. Compare the inventory with installed model files, configuration and observed network traffic.
Passed if	Inventory complete and consistent.

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.
Evidence	The vendor's terms and, if used, the third-party AI's terms; default settings.
Audit	1. Review the terms. 2. Check the product's default settings.
Passed if	No training without explicit opt-in; opt-in off by default.

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.

Evidence	Description of all telemetry and diagnostic data sent.
Audit	1. Process test documents containing canary values. 2. Capture all outgoing traffic, including crash reports (trigger a crash where possible). 3. Search the capture for canary values.
Passed if	No canary value outside the transfers the tier allows.

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.
Evidence	Signed label agreement.
Audit	1. Check the label agreement. 2. At spot checks and renewal: compare the architecture with the last audit report.
Passed if	Agreement signed; no unreported changes.

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.
Evidence	Connection list (table in the self-assessment).
Audit	1. Run the product behind the audit proxy with packet capture. 2. Process the test corpus. 3. Compare observed connections and DNS queries with the list.
Passed if	No connection that is not on the list.

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.
Evidence	Data processing agreement template, list of sub-processors.
Audit	1. Review the agreement and the list. 2. Compare the list with the model inventory and the connection list.
Passed if	Agreement available; list complete. Or: the vendor receives no customer data.

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.

Threshold	**Replaced: ≥ 99 % of structured identifiers (IBAN, e-mail, phone, ID and tax numbers), ≥ 98 % of person names and ≥ 95 % of all canary values – as B-03.**
Evidence	Description of the report function, screenshots of the preview, purpose stated in the terms.
Audit	1. Call the report function with texts of the test corpus. 2. Capture the reports sent, search them for canary values and compute the rates per group. 3. Check that nothing is sent without the user triggering it and that the preview matches what is sent. 4. Review the terms on how reports are used.
Passed if	All four conditions met and all three thresholds reached – or the product sends no reports.

4.2 Gold

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.
Evidence	Architecture description; location of the model files.
Audit	1. Locate the installed model files. 2. Process the test corpus and capture the traffic. 3. Search the capture for canary values and AI requests.
Passed if	No AI request and no canary value in outgoing traffic.

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.
Threshold	**Continued offline operation ≥ 14 days.**
Evidence	Description of licence check and updates.
Audit	1. Disconnect the network. 2. Process the test corpus with all AI functions. 3. Move the system clock forward 14 days, or verify the vendor's stated grace period.
Passed if	All AI functions work offline; grace period ≥ 14 days.

G-03 Verifiable models

Requirement	Models and model updates are delivered as files with a published checksum or signature.
Evidence	Where the checksums or signatures are published.
Audit	1. Recompute the checksum or verify the signature of a delivered model file.
Passed if	Checksum or signature matches.

4.3 Silver

S-01 Vendor-operated AI (Variant S-a)

Requirement	AI that processes personal data runs at the customer's site or on vendor-operated infrastructure in the EEA. No third-party AI sees personal data. A cloud provider's model services count as third-party AI.
Evidence	Architecture description, hosting contracts, locations.
Audit	1. Review hosting contracts and locations. 2. Obtain proof that the vendor runs the models itself (configuration, deployment, model files). 3. Process the test corpus through the service; check the vendor servers' outgoing traffic where accessible, otherwise the vendor's logs.
Passed if	Models demonstrably run at the vendor in the EEA; no third-party AI receives canary values.

S-02 Router (Variant S-b)

Requirement	If a third-party AI is used, a router at the customer's site or on vendor-operated infrastructure checks every request. Only requests without personal data go to the third-party AI.
Evidence	Description of the router (method, location, categories).
Audit	1. Trace the router in the architecture. 2. Check that no request reaches the third-party AI bypassing the router.
Passed if	Every request to the third-party AI passes the router.

S-03 Router accuracy (Variant S-b)

Requirement	The router reliably keeps requests containing personal data away from the third-party AI.
Threshold	**Of the test requests containing personal data, at most 0.5 % reach the third-party AI (≥ 99.5 % kept back); for requests containing structured identifiers (IBAN, e-mail, phone, ID and tax numbers), at most 0.1 % (≥ 99.9 %).**
Evidence	The vendor's own measurements (optional).
Audit	1. Send at least 1,000 test requests per language (mixed, with and without personal data). 2. Count which requests containing canary values reach the third-party AI. 3. Compute the rates separately for structured identifiers and other data.
Passed if	Both thresholds met, in German and in English.

S-04 Router fails closed (Variant S-b)

Requirement	If the router is uncertain, fails or cannot be reached, the request stays local or with the vendor.
Evidence	Description of the failure behaviour.
Audit	1. Deliberately disturb the router (stop it, force a timeout, return an invalid result). 2. Check whether any request reaches the third-party AI.
Passed if	Under no disturbance does a request reach the third-party AI.

4.4 Bronze

B-01 Pseudonymisation before transfer

Requirement	Every request to a third-party AI passes a pseudonymisation step at the customer's site or on vendor-operated infrastructure.
Evidence	Architecture description of the pseudonymisation step.
Audit	1. Trace the pseudonymisation step in the architecture. 2. Check that no request reaches the third-party AI bypassing it.
Passed if	Every request to the third-party AI is pseudonymised first.

B-02 Categories covered

Requirement	At least the following are covered: names of natural persons, addresses, dates of birth, phone numbers, e-mail addresses, IBAN and account numbers, ID card, passport, tax, social security and health insurance numbers, file and case reference numbers, vehicle registration numbers, IP addresses. The vendor lists any further categories.
Evidence	List of categories covered.
Audit	1. Review the list. 2. Process canary values of each category in the test corpus (see B-03).
Passed if	All minimum categories covered.

B-03 Detection rate

Requirement	The pseudonymisation step reliably removes personal data. A canary value counts as leaked if it, or an obvious variant of its spelling, appears in plain text in the outgoing request.
Threshold	**Removed: ≥ 99 % of structured identifiers (IBAN, e-mail, phone, ID and tax numbers), ≥ 98 % of person names and ≥ 95 % of all canary values overall.**
Evidence	The vendor's own measurements (optional).
Audit	1. Process at least 200 test documents per language. 2. Capture the outgoing requests to the third-party AI. 3. Determine for each canary value whether it appears in plain text; compute the rates per group.
Passed if	All three thresholds met, in German and in English.

B-04 Mapping table stays

Requirement	The mapping table (placeholder ↔ real value) never leaves the customer or the vendor. Placeholders in AI output are replaced back at the customer's site or at the vendor.
Evidence	Description of where the mapping table is stored and when it is deleted.
Audit	1. Search outgoing requests for mappings (placeholder next to canary value).
Passed if	No mapping in outgoing traffic.

B-05 Pseudonymisation fails closed

Requirement	If pseudonymisation fails or cannot be reached, nothing is sent to the third-party AI.
Evidence	Description of the failure behaviour.

Audit	1. Deliberately disturb the pseudonymisation (stop it, force a timeout, return an invalid result). 2. Check whether any request reaches the third-party AI.
Passed if	Under no disturbance does a request reach the third-party AI.

B-06 Visible to users

Requirement	Users can see what was sent to the third-party AI (preview before sending, or a log).
Evidence	Screenshots of the preview or log.
Audit	1. Open the preview or log in the product and compare it with the capture.
Passed if	Preview or log available and complete.

B-07 Contract with the third-party AI

Requirement	The contract with the third-party AI provider excludes training on the data, limits retention to at most 30 days, and provides for processing in the EEA or a valid basis for transfers to third countries.
Threshold	**Retention ≤ 30 days.**
Evidence	Contract or terms of the third-party AI, selected region.
Audit	1. Review the contract and configuration (region, retention).
Passed if	No training, retention ≤ 30 days, valid transfer basis.

B-08 Correct wording

Requirement	The product and its marketing do not call pseudonymised data anonymous or anonymised.
Evidence	Addresses of product and marketing texts.
Audit	1. Review product and marketing texts.
Passed if	Nothing described as anonymous.

5 NOCIP index

Within the tier, the NOCIP index rates the implementation with 0 to 10 whole points: five groups of 0 to 2 points each. The index never changes the tier; it is shown next to it (e.g. “Silver · 8/10”).

Group	1 point	2 points
I-1 Transparency	A data-flow diagram is published that shows the processing location of every AI function.	In addition, the product shows for each AI function or request where it is processed.
I-2 Customer control	The customer can switch off every external AI connection; the product then continues with reduced functions. (Gold: met.)	In addition, all AI functions are available in a variant that runs entirely at the customer's site. (Gold: met.)
I-3 Data minimisation	AI input and output are kept outside the customer for at most 30 days. (Gold: met.)	In addition, nothing is kept outside the customer beyond processing, and local logs contain content only if the customer switches this on.

Group	1 point	2 points
I-4 Integrity	The product verifies the signature or checksum of every model before loading it.	In addition, local stores holding content or mapping tables are encrypted (or set-up with operating-system encryption is documented).
I-5 Verifiability	The customer can view a log of all external AI requests (Gold: of all outgoing connections).	In addition, the vendor gives NOCIP permanent test access or a test installation for unannounced spot checks.

6 Rules

- **Tier:** A product receives the highest tier whose criteria it fully meets; the common criteria (C) always apply. Silver is met when all criteria of the applicable variant(s) are met.
- **Validity:** A label is valid for 12 months for the audited product version and later versions without changes to the AI architecture. Renewal requires a new audit.
- **Register:** Every label receives an ID (e.g. NOCIP-G-2026-0001) and a public entry on nocip.com with product, version, tier, index, catalogue version and validity. The label's seal carries the ID as a number in its ring (e.g. “GEPRÜFT · NR. G-2026-0001”).
- **Suspension and withdrawal:** If a change leads to a lower tier, the label is suspended until it is re-audited. False statements or misuse lead to withdrawal; the withdrawal stays visible in the register.
- **Fees:** The audit and the label are free of charge. One retest after a failed criterion is free of charge.
- **Auditor:** Audits are carried out by be next GmbH. be next itself offers pseudonymisation software (LexITanon); it audits its own products against the same checklist and publishes the full audit report.

7 Audit environment

- Test environment: the product runs on a NOCIP test machine or as test access. All traffic passes through an intercepting proxy (e.g. mitmproxy) with packet capture (all protocols, not only HTTP) and a DNS log.
- TLS: the audit certificate is installed. If certificate pinning prevents inspection, the vendor provides a test build without pinning or a log of outgoing data; otherwise the affected criteria count as not demonstrated.
- Canary values: new, unique canary values are generated for every audit, so that every match can be traced.
- Test corpus: the current version of the NOCIP test corpus (German and English) is used; the version is stated in the audit report.
- Retention: captures, logs and screenshots are kept for 3 years and not passed on.