

Data Protection Impact Assessment (DPIA)

CIATA DeepResearch — AI-Powered Web Research Engine

Version: 1.8 **Controller:** CIATA GmbH, Arnimallee 7, 14195 Berlin, Germany HRB 266084 B, Amtsgericht Charlottenburg Managing Director: Prof. Dr. Tim Landgraf **Privacy Contact / DPO:** privacy@ciata.io **Assessment Date:** 2026-06-18 **Last Revised:** 2026-09-27 **Next Review:** 2027-06-18

NOTE: This document has not yet undergone external legal review. It reflects the technical state of the system as at 2026-09-10. See Section 7 (DPO Consultation) before relying on this assessment for regulatory purposes.

Changelog

Version	Date	Summary
1.0	2026-06-18	Initial assessment
1.1	2026-07-10	Sub-processor table completed (all 8 entries with data-sent column); risk methodology formalised (3×3 L×S matrix with Recital 75 harm types); missing risks added (account-data breach, Art. 9 research queries, EEA–US backend transfer, Staan full_content relay); TOM corrections (self-hosted Plausible analytics; VSRMark watermark implemented; C2PA remains planned; Google Fonts CDN risk resolved); Necessity & Proportionality expanded into substantive paragraphs; DPO consultation section added; residual-risk rating and Art. 36 reasoning added; Art. 6(1)(c) AI Act legal-basis row added; DSAR/erasure handling note added; DPIA trigger subsection added
1.8	2026-09-27	Consistency Checker added (issue #172); runs in parallel with the Paragraph Checker on qwen3.8-27b; no new model, processor or data category; two additional LLM calls per report. Sub-processor table updated: Consistency Checker added to AKI.io qwen3.8-27b agent list.
1.7	2026-09-27	Language Guard moves from qwen3.6-chat-35b to gemma4-chat-26b with reasoning disabled (issue #167). qwen3.6-chat-35b is no longer used in any agent; three models remain in use. Failed repairs now remove leaked characters in code. Sub-processor table updated: AKI.io language guard row changed from qwen3.6-chat-35b to gemma4-chat-26b. Same data categories, same recipient.
1.6	2026-09-26	Transactional email provider changed from Brevo SAS to Scaleway SAS (issue #152); Scaleway standard DPA (version June 2024) pending acceptance and filing. Report-ready notification mail added (issue #159): email address, name, report title, short summary (up to ~500 characters), link, deletion date now sent to Scaleway. Sub-processor table updated: Scaleway SAS replaces Brevo SAS (former

Version	Date	Summary
		provider, DPA never executed, wind-down pending). New risk R10 added for research content transmitted by report-ready email.
1.5	2026-09-13	Model routing revised for run time (issue #135): the Relevance Gate moves from qwen3.6-chat-35b to gpt-oss-120b (same verdicts on 192 replayed results, 4.5x faster) and the Extractor from qwen3.6-chat-35b to gemma4-chat-26b with reasoning disabled (5x faster, more and better-grounded findings on 40 real pages). gemma4-chat-26b is a NEW model identifier — a fourth model now processes customer data (PII-filtered page excerpts) — so the DPA scope must name four models; qwen3.6-chat-35b remains in service for the Language Guard only. Same data categories, same recipient (AKI.io GmbH, Frankfurt).
1.4	2026-09-12	Post-synthesis verification restructured (issue #133): the Fact Checker — a whole-report check plus a per-citation alignment check, both on gpt-oss-120b — is replaced by the Paragraph Checker, and the Report Patcher is given its own routing; both now run on qwen3.8-27b. Same data categories (the draft report and the run's findings), same recipient (AKI.io GmbH), no new model identifier: the three-model set and the DPA scope instruction are unchanged.
1.3	2026-09-10	Synthesis model corrected from mistral4-119b to qwen3.8-27b, in service since 2026-08-23 (commit eb2a5b1); the Table Repair agent, added 2026-09-10, runs on that same model and introduces no new recipient. DPA scope instructions corrected in step. Extractor token budget raised 2,560 → 5,120; the Language Guard borrows that configuration, so its per-line ceiling moved with it. No change to what data is sent, to whom, or for how long it is kept.
1.2	2026-08-11	Corrected page-content storage description (PII-filtered text IS stored in state.json for 30-day TTL, readable only by a user who holds the admin role AND owns that run, with no administrative override; "transient" label removed); added three undisclosed data categories (company_profiles B2B billing table v8, withdrawal_consent §356(5) BGB consent table v10, run provenance metadata); added TOM for server-side JWT token revocation on password change/reset (schema v12 token_version); verified no remaining MiniMax model references; updated technical-state date throughout

1. Processing Description

1.1 DPIA Trigger and Necessity Assessment

A DPIA is required under **Art. 35(1) GDPR** where processing "is likely to result in a **high risk** to the rights and freedoms of natural persons." Two paths can establish that: the three automatic cases listed in Art. 35(3), and the general high-risk assessment under Art. 35(1) informed by the nine EDPB criteria (WP248 rev.01 — meeting two or more indicates a DPIA is likely required). **This assessment concludes that a DPIA is required primarily via the Art. 35(1) / WP248 path.** The Art. 35(3) automatic cases are examined below but are not relied upon as clearly met — this is a reasoned high-risk determination, not a black-letter Art. 35(3) case.

Art. 35(3) — automatic cases (examined; not clearly met):

Art. 35(3) case	Assessment
(a) Systematic, extensive profiling producing legal or similarly significant effects	Not applicable. The service does not evaluate, score, or make decisions about the individuals whose data it encounters. No automated decision-making within the meaning of Art. 22 GDPR takes place.
(b) Large-scale processing of special-category (Art. 9) data	Not clearly met. CIATA does not target or intentionally process Art. 9 data. Such data may appear incidentally in scraped content, but incidental capture is not "large-scale processing of special categories" as a core activity.
(c) Systematic monitoring of a publicly accessible area on a large scale	Not clearly met. This limb (Recital 91) is directed at behavioural monitoring of individuals in observed spaces (e.g. CCTV). Applying it to scraping of public web pages is an analogy, not a clear fit, and is not relied upon here.

Art. 35(1) – high-risk assessment (the operative basis). Against the WP248 rev.01 criteria, CIATA's processing engages:

WP248 criterion	Applies?	Basis
Innovative use of a new technology (GPAI / LLMs)	Yes – strong	The engine relies on generative large language models to autonomously decompose, search, extract and synthesise content – a novel processing form expressly identified by regulators as risk-relevant (WP248 criterion 8).
Systematic processing of data on subjects with no relationship to the controller	Yes – moderate	Individuals whose personal data appears on scraped pages have no relationship with CIATA and limited practical ability to object (see §3 Risk Assessment and the companion LIA).
Data processed on a large scale	Arguable	High volume of web content is processed, but the personal data per page is typically incidental and thin. Whether this constitutes "large-scale processing of personal data" (WP248 factors: number of subjects, volume, duration, geographic extent) is a genuine judgment.
Sensitive / highly personal data	Possible – weak	Art. 9 data may appear incidentally; it is neither sought nor the purpose.
Data concerning potentially vulnerable persons	Possible – weak/moderate	Depending on the query, content may concern consumers, patients, or others in vulnerable positions.

The innovative-technology criterion, combined with systematic processing of unrelated third parties' personal data obtained by web scraping, places this processing over the WP248 "two-or-more criteria" threshold at which a DPIA is indicated. **This DPIA is therefore conducted on the basis that the processing is likely to result in a high risk under Art. 35(1) – primarily by reason of the combination of novel generative-AI technology and systematic processing of third-party personal data.** Conducting it is in any event prudent and is treated here as required.

△ **TO CONFIRM (external legal review):** Because this rests on an Art. 35(1) high-risk judgment rather than a clear Art. 35(3) trigger, two points need confirmation: (i) whether

CIATA's processing appears on the **Berliner Beauftragte für Datenschutz und Informationsfreiheit (BlnBDI) list of processing operations requiring a DPIA** (Art. 35(4)); and (ii) counsel's view on the high-risk determination. This is consequential beyond this document: under **§ 38(1) sentence 2 BDSG**, processing that is subject to a mandatory DPIA also triggers a **mandatory Data Protection Officer regardless of headcount** — so the DPIA-necessity and DPO-necessity questions (see the DPO Consultation section) stand or fall together.

1.2 Nature of Processing

CIATA DeepResearch is an autonomous research engine that:

1. Accepts a natural-language research question from an authenticated user
2. Decomposes the question into sub-questions using a GPAI language model
3. Generates search queries and executes them via a privacy-preserving search API (Staan, default backend)
4. Fetches and reads publicly accessible web pages — either directly (with robots.txt compliance and per-domain rate limiting) or via Staans server-side full_content mode, which retrieves page text on CIATA's behalf
5. Applies a regex-based PII filter to all scraped content before storage to job state and before any AI processing (the filtered text is the only form written to runs/{job_id}/state.json and the only form sent to the extractor LLM)
6. Extracts factual findings from PII-filtered content using AI
7. Synthesizes findings into a structured research report
8. Embeds an invisible VSRMark Unicode-variation-selector watermark and a visible "AI-generated" label and provenance HTML comment in every report (EU AI Act Art. 50(2))
9. Delivers the watermarked report to the user via the API

1.3 Scope

- **Data Subjects:** Authenticated users (account holders); incidentally, individuals mentioned in publicly accessible web page content that is scraped and processed
- **Data Categories:**
 - User account data (email address, bcrypt-hashed password)
 - Research queries and AI-generated research reports
 - Scraped web page content (PII-filtered before both storage and LLM submission; capped at page_max_chars, 16,000 characters in production; stored as page_text within runs/{job_id}/state.json for the 30-day job TTL; accessible via GET /v1/admin/reports/{job_id}/sources/{source_index}/text only to a user who simultaneously holds the admin role AND owns the job — there is no staff-level override that allows reading another user's stored page text, see _require_admin() and _check_job_ownership() in src/api/app.py)
 - Company billing profiles (company_profiles table, schema v8): user_id, organization_name, given_name, family_name, email, street_and_number, postal_code, city, country, vat_number, registration_number. B2B customers only; stored in CIATA's

SQLite database on Hetzner DE; retained until account deletion (§147 AO interaction for invoice-related fields requires counsel confirmation — see action A8)

- Withdrawal consent records (withdrawal_consent table, schema v10): job_id, user_id, consent_version, locale, consented_at. Evidence of a consumer's §356(5) BGB immediate-performance consent; stored in CIATA's SQLite database on Hetzner DE; NOT subject to the 30-day job-state TTL; minimum retention period to be confirmed by counsel (statutory limitation periods, §§195/199 BGB — see action A7)
- Payment data (amount, pack reference, user identifier; card details processed exclusively by Mollie B.V.)
- Billing address (collected only for B2B/Billie payment method users)
- Run provenance metadata (schema_version, code_revision, agent_models, config_digest; serialised into every state.json): operational metadata, not personal data; 30-day TTL with job state
- Technical/server-side access logs
- **Processing Volume:** Approximately 100–1,000 research jobs per month (estimated)
- **Geographic Scope:** EEA by default (all primary infrastructure in Germany/EU); EEA-to-US transfer is possible if a non-default search backend is configured (see Risk R6)

1.4 Purpose

To provide B2B customers with AI-powered deep research capabilities for market analysis, competitive intelligence, regulatory research, and academic literature review.

1.5 Legal Basis

Data Type	Legal Basis	Justification
Account data (email, hashed password)	Art. 6(1)(b) — Contract	Necessary to provide the service and manage authenticated user accounts
Research queries and generated reports	Art. 6(1)(b) — Contract	User-initiated, asynchronous processing of the core contractual deliverable
Scraped web page content (PII-filtered; stored for 30-day TTL)	Art. 6(1)(f) — Legitimate Interest	Necessary for B2B research functionality using publicly accessible information; data minimisation via PII filter applied before storage and before LLM submission; contractual prohibition on AI training use. Separate Legitimate Interest Assessment (LIA) on file.
Company billing profiles (company_profiles table)	Art. 6(1)(b) — Contract + Art. 6(1)(c) — Legal obligation	(b) Necessary to perform the B2B sales contract (invoicing, Qonto invoice generation, Billie credit underwriting); (c) §147 AO requires retention of business records including invoices. ⚠ Interaction between Art. 6(1)(b) account-deletion erasure right and the §147 AO minimum retention obligation requires counsel confirmation (see action A8).
Withdrawal consent records (withdrawal_consent table)	Art. 6(1)(b) — Contract	Documentary evidence that the consumer expressly requested immediate performance and acknowledged that the right of withdrawal lapses on

Data Type	Legal Basis	Justification
		generation (§356(5) / §312f BGB). Deletion on account erasure would destroy legal evidence – minimum retention to be confirmed by counsel (§§195/199 BGB limitation periods; see action A7).
Payment data	Art. 6(1)(b) – Contract	Necessary for credit purchase processing via Mollie
Technical/access logs	Art. 6(1)(f) – Legitimate Interest	Security monitoring, debugging, and fraud prevention; proportionate to operational necessity
AI-generated content disclosure metadata (VSRMark watermark, visible label, provenance comment)	Art. 6(1)(c) – Legal obligation	EU AI Act Art. 50(2) requires deployers of GPAI systems to ensure that AI-generated content is marked in a machine-detectable and human-readable format

2. Necessity and Proportionality

2.1 Necessity Assessment

Research queries and reports are strictly necessary to provide the core contractual service. The asynchronous research pipeline cannot function without storing the query (to track job state) and the generated report (for user retrieval). No less intrusive technical design achieves the same outcome.

Web scraping is the primary data acquisition method and is operationally necessary: AI research synthesis requires access to primary source text. Alternatives – such as restricting the service to licensed database content – would fundamentally change the product offering and fall outside CIATA's current business model. Direct fetching is supplemented by Staan's full_content mode (see §2.3) to overcome bot-wall restrictions that commonly result in 403 errors on publisher sites.

Account data (email, hashed password) is the minimum necessary for authentication and authorisation. CIATA does not collect name, telephone number, or address at registration. B2B users who elect to use Billie invoice payment provide billing address data at checkout, which is necessary for Billie's credit underwriting.

Payment data: Card details are processed exclusively by Mollie B.V. and never transmitted to or stored by CIATA. CIATA retains only Mollie metadata (pack_id, credits, user_id, Mollie payment reference) necessary for credit granting, idempotency, and Qonto invoice generation.

2.2 Data Minimisation Measures

- **PII filtering (regex-based):** Applied to all scraped web content before storage to disk and before submission to any LLM. The scrub (scrub_pii() in src/pii_filter.py) runs at

src/engine.py line 859 — before `source.page_text` is assigned — so what is written to `runs/{job_id}/state.json` and what is sent to the extractor are the same already-filtered text. The filter covers: email addresses, international telephone numbers (including German/Austrian/Swiss formats), DE Steuerliche Identifikationsnummer (11-digit), DE Sozialversicherungsnummer, Swiss AHV/AVS numbers, IBANs, Luhn-validated credit card numbers, IPv4 addresses, and German-format street addresses (e.g., "Arnimallee 7"). False-positive guards prevent redaction of years, version numbers, and other legitimate numeric sequences. Known limitations (personal names and non-DACH national identifiers not covered) apply to both the stored copy and the LLM submission — see §2.3 and the R1 residual.

- **robots.txt compliance:** The page reader (`src/page_reader.py`) checks `robots.txt` before every direct URL fetch, using a per-domain in-memory cache with a 1-hour TTL. Where `robots.txt` is unreachable (connection error, timeout, HTTP 404), the URL is allowed by graceful fallback — a proportionate design to avoid blocking legitimate public-interest content without the consent of the site.
- **Honest bot identification:** The User-Agent string identifies the agent as `CIATA-DeepResearch/1.0` and includes the privacy contact (`privacy@ciata.io`), consistent with EU AI Act transparency norms and EDPB guidance on web scraping.
- **Content truncation:** Directly fetched pages are truncated to a maximum of 16,000 characters (production configuration via `config.yaml` override; code default is 8,000 as defined in `ResearchDefaults.page_max_chars`).
- **Self-hosted assets:** Google Fonts CDN and `marked.js` were migrated to CIATA-controlled self-hosting on 2026-07-10, eliminating the prior incidental EEA-to-US data transfer that occurred on each page load (resolved risk — see §4.1).
- **Data retention:** Research reports and job state (including stored PII-filtered `page_text`): 30 days (`DATA_RETENTION_DAYS`, enforced by automated filesystem cleanup in `src/api/job_store.py` — deletes run directories only, not SQLite rows). Server/access logs: approximately 90 days. Payment records: 10 years (§ 147 AO tax retention obligation). JWT refresh tokens: 7 days. Unverified account registrations: 72 hours. Company billing profiles (`company_profiles` table): retained until account deletion; the §147 AO minimum retention obligation for invoice-related data fields requires counsel confirmation before an erasure request can be fully honoured on those fields (see action A8). Withdrawal consent records (`withdrawal_consent` table): NOT deleted with the 30-day job-state TTL (the filesystem cleanup does not touch SQLite rows); NOT deleted on account soft-deletion (`soft_delete_user` does not delete consent rows, which is correct — they are legal evidence); minimum retention period requires counsel confirmation (§§195/199 BGB limitation periods — see action A7).
- **No AI training:** User data is contractually prohibited from use in AI model training under CIATA's DPA with AKI.io.

2.3 Proportionality

30-day retention: The 30-day report TTL balances user utility (accessing research results after initial delivery) against storage minimisation. Research reports may contain incidental personal data about third parties embedded in quoted web content. A 30-day window is

proportionate because: (a) B2B users typically retrieve results within hours to days of completion; (b) automatic TTL-based deletion eliminates the need for user-initiated erasure requests in the vast majority of cases; (c) longer retention would increase the risk surface without a corresponding benefit to the contractual purpose of the service.

Regex PII filter vs. LLM-based Named Entity Recognition (NER): The regex approach was deliberately chosen over LLM-based NER. Sending raw scraped text to a second LLM call for PII detection would introduce additional data transfers, increase processing cost, and add latency – without eliminating residual risk, since NER also has meaningful false-negative rates on contextual personal data. Regex minimises data transfers and is fully deterministic and auditable. Known limitations – the filter does not redact personal names, non-DACH national identity numbers (UK NI, US SSN, FR/NL/PL equivalents), or addresses in non-German formats – are accepted as proportionate under the Legitimate Interest basis, on the grounds that: (a) the content is already publicly accessible; (b) AKI.io's no-training commitment is intended to prevent this data from influencing model training – noting that the DPA giving it effect is not yet executed (§6); and (c) a commercially focused DACH-region research service has materially lower exposure to non-DACH PII formats.

Staan full_content mode: When the Staas search backend is used with full_content=true (the production default), Staas's servers fetch each result page server-side and return the full text as markdown. Page content therefore passes through Staas's EU-based infrastructure before reaching CIATA's own PII filter. This design is proportionate because: (a) it overcomes bot-wall 403 errors that a direct CIATA crawler would encounter on many publisher sites, which is necessary for research quality; (b) Staas operates within the EU and is engaged as a processor – note the Art. 28 DPA with Staas is **required but not yet executed** (see §6), so this limb rests on intended contractual terms rather than a concluded agreement; (c) CIATA applies its own PII filter to Staas's output before storage to disk and before any submission to the LLM, maintaining the data minimisation chain through both the stored copy and the LLM path. The alternative – direct fetching by CIATA – would produce lower-quality results without reducing the aggregate data surface for data subjects.

3. Risk Assessment

3.1 Risk Methodology

Risks are evaluated on a 3×3 Likelihood × Severity matrix, following EDPB WP248 rev.01. Harm types are drawn from Recital 75 GDPR: discrimination; identity theft or fraud; financial loss; damage to reputation; loss of confidentiality of professionally protected data; unauthorised reversal of pseudonymisation; and significant economic or social disadvantage.

Likelihood Scale:

Level	Score	Definition
Low	1	

Level	Score	Definition
		Unlikely; occurrence requires extraordinary circumstances or multiple simultaneous control failures
Medium	2	Possible; might occur without extraordinary circumstances under normal operations
High	3	Probable; likely to occur routinely under normal operating conditions

Severity Scale (Recital 75 harm types):

Level	Score	Definition
Low	1	Minimal impact; data subjects unlikely to experience any discernible negative consequence from the Recital 75 harm types
Medium	2	Significant impact; data subjects may experience financial loss, reputational damage, discrimination, or loss of confidentiality
High	3	Serious or irreversible impact; includes identity theft, substantial financial harm, discrimination on Art. 9 grounds, or severe social disadvantage

Risk Rating Thresholds:

Score (L × S)	Rating
1–2	Low
3–4	Medium
6–9	High

Pre-mitigation and post-mitigation (residual) ratings are both recorded. A residual rating of High would trigger Art. 36 prior consultation.

3.2 Risk Register

#	Risk	Pre-mitig. L	Pre-mitig. S	Pre-mitig. Score	Controls / Mitigations	Residual Rating
R1	Incidental structured PII (email, phone, ID numbers, IBAN, card, IPv4, DE address) in scraped content stored in job state and/or sent to LLM via Staan full_content relay or direct fetch	2	2	4 — Medium	Regex PII filter applied to all web content before storage to disk and before LLM submission (src/pii_filter.py; applied at engine.py:859 before source.page_text is assigned); 30-day TTL deletes stored page_text with job state; no-training commitment from AKL.io and coverage of the Staan relay step are the intended contractual basis — neither DPA is	Low-Medium — personal names and non-DACH identifiers not covered by regex remain in both the stored copy and the LLM submission (see §2.3); residual accepted under LIA

#	Risk	Pre-mitig. L	Pre-mitig. S	Pre-mitig. Score	Controls / Mitigations	Residual Rating
					executed (see §6); stored page_text readable only by the job-owning user who also holds the admin role (no staff override)	
R2	Unauthorised access to a user's research reports by another authenticated user or unauthenticated actor	1	2	2 — Low	JWT authentication required on all data endpoints; per-endpoint ownership verification on all GET /v1/research/{job_id} responses; HTTPS/TLS via Caddy	Low
R3	Server-side data breach of stored research reports and job state	1	2	2 — Low	HTTPS/TLS in transit; bcrypt-hashed passwords; Hetzner infrastructure physical and logical security; strict per-user application-level access isolation; automated 30-day retention cleanup reduces the live data set; encrypted off-site backups held geo-redundantly within the EU (Frankfurt → Helsinki)	Low — production disks are not encrypted at rest (accepted gap with compensating controls; see below)
R4	Account-data breach (user email addresses and bcrypt password hashes)	1	2	2 — Low	bcrypt hashing with appropriate work factor; rate limiting on authentication endpoints; JWT expiry (60 min access, 7 day refresh); unverified accounts auto-deleted at 72 hours	Low
R5	Research question or report contains personal data or Art. 9 special-category data about a third party (e.g., query: "investigate [person X]'s health status")	2	2	4 — Medium	30-day retention TTL eliminates data over time; AKI.io no-training commitment (intended contractual basis; DPA not yet executed); B2B ToS prohibits use for discriminatory profiling; B2B customers are upstream data controllers responsible for query subject matter	Low-Medium — CIATA cannot comprehensively screen incoming queries; residual accepted via contractual framework

#	Risk	Pre-mitig. L	Pre-mitig. S	Pre-mitig. Score	Controls / Mitigations	Residual Rating
R6	EEA-to-US data transfer when a non-default search backend (DuckDuckGo US servers, Google Custom Search) is activated	1	2	2 — Low (default config)	Staan (EU-sovereign, Germany/EU infrastructure) is the configured production default (backend: <code>staan</code> in <code>config.py</code>); US backends exist in <code>src/searcher.py</code> but require explicit operator configuration via <code>SEARCH_BACKEND</code> env var or <code>config.yaml</code> ; no SCCs in place for DuckDuckGo or Google CSE	Low in default production — conditional on prod configuration (△ see below)
R7	Publicly accessible web page text relayed by Staas in <code>full_content</code> mode passing through a third-party EU processor before CIATA's PII filter	3	1	3 — Medium	DPA with Staas; EU-sovereign Staas infrastructure (Ecosia + Qwant, EU-based); CIATA PII filter applied to all Staas output before any LLM submission; content is already publicly accessible	Low — Staas is EU-based under DPA; relay does not meaningfully increase risk to data subjects relative to a direct CIATA fetch
R8	LLM hallucination producing false statements about identifiable natural persons in a research report	1	1	1 — Low	Visible "AI-generated" label on every report; VSRMark invisible watermark for machine-detectable provenance (EU AI Act Art. 50(2)); user verification recommendation included in all reports	Low
R9	EEA-to-third-country transfer via primary LLM API	N/A	N/A	N/A	AKI.io processes exclusively in Frankfurt, DE; GDPR applies directly; no adequacy decision needed; no international transfer occurs on the default LLM path	N/A
R10	Research content (report title, short summary of up to approximately 500 characters) transmitted to the	2	2	4 — Medium	Email processor limited to the EU (fr-par, Paris, France; Scaleway TEM documentation states EU-only hosting and processing); no open or	Low

#	Risk	Pre-mitig. L	Pre-mitig. S	Pre-mitig. Score	Controls / Mitigations	Residual Rating
	user's inbox via the transactional email processor. The summary could reach unintended third parties if the email is forwarded. Relevant for professional-secrecy holders (§ 203 StGB, e.g. law firms): the report subject could be disclosed via the notification.				click tracking in Scaleway TEM (links delivered as written, no tracking-domain rewrite); summary capped at approximately 500 characters; report-ready notification can be turned off in account settings (Notifications); the full report remains on CIATA's servers, accessible only by the authenticated account holder; Scaleway DPA pending acceptance and filing	

✓ **CONFIRMED (2026-07-15, R3):** The Hetzner Cloud production instance does **not** use disk encryption at rest (no Volume-level or block-device/LUKS encryption). This is a documented and accepted gap. Compensating controls: bcrypt-hashed passwords, TLS in transit, strict per-user application-level access isolation, automated 30-day retention cleanup, and **encrypted off-site backups stored geo-redundantly within the EU (Hetzner Frankfurt → Helsinki)**. Recommendation: evaluate enabling volume-level (LUKS/dm-crypt) at-rest encryption before the hot-production cutover (see action A1).

⚠ **TO CONFIRM (R6):** Confirm that the production deployment sets `SEARCH_BACKEND=staan` (or equivalent) in the production `.env` or `config.yaml`, so that DuckDuckGo and Google CSE are not inadvertently reachable. If non-EU backends are ever used in production, Standard Contractual Clauses (SCCs — Implementing Decision 2021/914) or an equivalent Art. 46 transfer mechanism must be executed before activation.

3.3 Residual Risks

After all mitigations, the following residual risks remain and are accepted:

- **R1 (Incidental PII — residual):** The regex filter does not catch personal names, non-DACH national identifiers (UK NI, US SSN, FR/NL/PL equivalents), or non-German postal addresses. This applies equally to the stored `page_text` in `state.json` and to what is submitted to the LLM — the stored copy may therefore contain names and non-DACH identifiers of third parties for up to 30 days. This residual is accepted as proportionate under the Legitimate Interest basis: the content is already publicly accessible; the 30-day TTL limits the storage exposure window; access to stored `page_text` is restricted to the job-owning user who also holds the admin role; AKI.io's no-training commitment is intended to prevent use in model training, though the DPA giving it effect is not yet

executed; and LLM-based NER would not eliminate the residual and would introduce additional processing risk. Rating remains Low-Medium.

- **R5 (Art. 9 research queries — residual):** CIATA cannot prevent users from submitting research queries that concern identifiable natural persons or Art. 9 special-category data. B2B customer ToS and the B2B data processing agreement place responsibility on the customer as the upstream data controller for the query's subject matter. CIATA acts as a processor for these purposes.
- **R6 (EEA-to-US transfer — conditional):** This risk is residual only if the production configuration is confirmed to pin the Staan backend (pending — see action item A2).
- **R3 (At-rest encryption — accepted gap):** Confirmed (2026-07-15) that production disks are not encrypted at rest. The residual is accepted on the basis of the compensating controls above — in particular encrypted EU-geo-redundant off-site backups, strict per-user access isolation, and short (30-day) retention. Enabling volume-level at-rest encryption before the hot-production cutover remains a recommended hardening step (action A1).

Aggregate residual risk: Low to Low-Medium. No risk registers at High after mitigation.

4. Technical and Organisational Measures

4.1 Technical Measures

- [x] HTTPS/TLS encryption in transit (Caddy reverse proxy)
- [x] Encrypted off-site backups, stored geo-redundantly within the EU (Hetzner Frankfurt → Helsinki)
- [] Disk encryption at rest — **not enabled** on the Hetzner production instance (confirmed 2026-07-15); accepted gap with compensating controls (encrypted EU backups, per-user access isolation, 30-day retention). Enabling volume-level (LUKS/dm-crypt) encryption is a recommended pre-launch hardening step (action A1)
- [x] PII filtering on all scraped web content before storage to disk and before LLM submission (regex-based; applied at `engine.py:859` before `source.page_text` is assigned; see `src/pii_filter.py` and §2.2)
- [x] robots.txt compliance with per-domain in-memory cache (1-hour TTL) and graceful-fallback on fetch failures (see `src/page_reader.py`)
- [x] Honest bot identification: User-Agent string identifies CIATA DeepResearch and includes privacy contact
- [x] Per-domain rate limiting (minimum 1 second between requests to the same domain)
- [x] JWT authentication (access token: 60 min TTL; refresh token: 7 day TTL)
- [x] Server-side JWT token revocation on password change or reset: every outstanding access and refresh token is immediately invalidated, independently of TTL, by incrementing a `token_version` counter (`users.token_version`, schema v12) in the database. Every authenticated request validates the JWT's embedded `tv` claim against the stored

counter via `token_version_matches()` in `src/api/auth_users.py`; a mismatch returns HTTP 401. Password change (PATCH `/v1/auth/me/password`) and password reset (POST `/v1/auth/reset-password`) both call `db.bump_token_version()` — see `src/api/auth_users.py` lines 585 and 675.

- [x] Password hashing (bcrypt)
- [x] Per-endpoint ownership verification on all data-returning API endpoints, with no administrative override: `_check_job_ownership()` returns a 404 for any request where the authenticated `user_id` does not match the job's owner — including requests from users holding the admin role. CIATA staff cannot read another user's stored job data through any API endpoint.
- [x] Automated data retention cleanup (30-day TTL, `DATA_RETENTION_DAYS`, `src/api/job_store.py`)
- [x] AI-generated content disclosure — visible "AI-generated" label and provenance HTML comment embedded in every research report (EU AI Act Art. 50(2))
- [x] VSRMark invisible Unicode-variation-selector watermark (machine-detectable; implemented in `src/watermark.py`; applied in `engine.py` and `job_store.py`) — EU AI Act Art. 50(2)
- [x] Self-hosted privacy-preserving analytics: Plausible Analytics instance at `trc.ciata.io`, on CIATA-controlled infrastructure; cookie-free; processes no personal data; NOT a third-party sub-processor
- [x] Google Fonts and `marked.js` self-hosted since 2026-07-10 — eliminates the prior incidental EEA-to-US data transfer on each page load (resolved risk)
- [] C2PA content provenance manifest (planned, post-launch)

4.2 Organisational Measures

- Privacy policy published at `/privacy`
 - DPO contactable at `privacy@ciata.io` — see Section 7 (DPO Consultation) regarding designation and independence
 - Data processing agreements with third-party processors; the payment recipients Mollie, Billie and Qonto act as independent controllers and are disclosed under Art. 13/14 (see Section 6)
 - Annual DPIA review scheduled (next: 2027-06-18)
 - ⚠ Incident-response / data-breach procedure (Art. 33/34) — **NOT YET DOCUMENTED** (open action item)
 - No-training contractual commitment in CIATA's DPA with AKI.io
-

5. Data Subject Rights

Right	Implementation
Access (Art. 15)	Authenticated users: GET <code>/v1/research/</code> lists all jobs; GET <code>/v1/research/{job_id}</code> returns the full report. Formal DSAR: <code>privacy@ciata.io</code> (response within 30 days per Art. 12(3) GDPR).

Right	Implementation
Rectification (Art. 16)	AI-generated research reports are not records "about" the user and cannot be rectified as personal data of the requesting user. Account data (email): updateable via account settings.
Erasure (Art. 17)	DELETE /v1/research/{job_id} deletes a specific job and its associated files immediately. Self-service account deletion: DELETE /v1/auth/me (requires password confirmation; clears PII and deactivates the account) — or by contacting privacy@ciata.io . In-flight erasure: the background worker re-checks account <code>is_active</code> before a job starts and at each research iteration, so a deleted/deactivated account's queued or in-progress jobs are aborted rather than completed and stored. The 30-day automated retention TTL means research data is auto-deleted in the ordinary course. Payment records are subject to the 10-year §147 AO retention obligation and cannot be erased within that period.
Restriction (Art. 18)	Contact DPO at privacy@ciata.io . Processing will be restricted to storage only pending resolution.
Portability (Art. 20)	Research reports are downloadable in Markdown format via API.
Objection (Art. 21)	Data subjects may object to processing under Art. 6(1)(f) (Legitimate Interest) by contacting privacy@ciata.io . CIATA will assess the objection and cease processing unless it can demonstrate compelling legitimate grounds overriding the data subject's interests.

Third-party data subjects: Individuals whose personal data appears incidentally in scraped web content (e.g., their name or contact details mentioned on a public web page that was retrieved as part of a research job) may exercise rights under Art. 15–22 by contacting privacy@ciata.io. Given that CIATA does not intentionally collect data about third parties, and that such data is embedded in research report text that is auto-deleted after 30 days, access and erasure requests from third parties may be difficult to fulfil with precision. Where technically feasible, CIATA will endeavour to locate and delete identifiable third-party references within stored reports.

6. Sub-Processors

Note on Plausible Analytics: CIATA operates its own Plausible Analytics instance at `trc.ciata.io` on CIATA-controlled infrastructure (Hetzner). It is cookie-free, processes no personal data, and is NOT a third-party sub-processor. It is not listed in the table below.

Note on OpenObserve telemetry (PLANNED — not yet enabled): An operational metrics exporter exists in `src/perf.py` but transmits nothing unless `OPENOBSERVE_API_KEY`, `OPENOBSERVE_ORG` and `OPENOBSERVE_STREAM` are set, which they are not. When enabled it would ship fixed metric names, numeric values, an opaque per-run `job_id`, and a raw exception string on failures — no question text, report text, page content, URLs, email addresses or IP addresses. Because the `job_id` is correlatable to a user through CIATA's own database, the records are treated as pseudonymous personal data. Whether OpenObserve becomes a sub-processor depends on an undecided hosting choice: self-hosted on CIATA infrastructure would place it alongside Plausible above; a hosted service would make it an Art. 28 processor

requiring an executed DPA. It is therefore not listed in the table below. See ZDR_STATEMENT.md §5.9 for the full description and the pre-enablement conditions, and action item A9.

Sub-Processor	Purpose	Location	Data Sent	DPA Status
AKI.io GmbH	LLM inference — reasoning agents and the relevance gate: gpt-oss-120b; report synthesis, table repair and the post-synthesis paragraph check, consistency check and patch: qwen3.8-27b; extraction and language guard (reasoning disabled): gemma4-chat-26b. qwen3.6-chat-35b was used for the language guard until 2026-09-27.	Frankfurt, DE	Research queries; PII-filtered scraped page content; all LLM prompts and completions	⚠ Required (no-training DPA) — not yet executed / not on file
European Perspective SAS (Staan)	Web search — DEFAULT backend (Ecosia + Qwant index); with full_content=true (production default), Staan server-side fetches each result page and returns full page text as markdown, which CIATA then PII-filters before any LLM use	EU	Search queries; with full_content=true: full page text (markdown) of all search results	⚠ Required — not yet executed / not on file
Hetzner Online GmbH	Cloud infrastructure hosting	Germany, DE	All application data (Hetzner Cloud — shared virtualised infrastructure, not dedicated physical servers)	⚠ Required (Hetzner Cloud DPA) — not yet executed / not on file
Mollie B.V.	Card payment processing	Amsterdam, NL	Payment amount and currency; CIATA-internal user_id (in Mollie payment metadata); billing address if provided by the user	⚠ Required (Mollie DPA) — not yet executed; Mollie also acts as a co-controller for PCI DSS purposes
Billie GmbH	B2B Net-30 BNPL (embedded within Mollie checkout as payment method billie); performs	Berlin, DE	Billing address (given name, family name, organisation name, street, postal code, city, country); VAT number	Independent controller (BNPL credit decision, reached via Mollie's checkout)

Sub-Processor	Purpose	Location	Data Sent	DPA Status
	credit underwriting for invoice payment		(USt-IdNr.) and Handelsregister number for B2B credit underwriting	— no Art. 28 DPA required; Art. 13/14 disclosure in privacy policy §6. Classification to confirm with counsel.
Qonto	Invoice generation — called for every completed purchase to auto-generate a finalised Qonto client invoice via the Qonto v2 API	Paris, FR	Name, email address, billing address (for all purchasers, sourced from the stored company profile); VAT number if provided	Independent controller (licensed payment institution acting in its own name) — no Art. 28 DPA required; Art. 13/14 disclosure in privacy policy §6. Classification to confirm with counsel.
Scaleway SAS	Transactional email (verification, password reset, purchase confirmation, report-ready notification)	Paris, FR (fr-par, EU-only)	Email address; name (if given); email verification and password reset links; purchase details (credits, pack label, price, payment reference, purchase date); report-ready notification: report title, short summary (up to ~500 characters), link to report, deletion date	⚠ Required — Scaleway standard DPA (version June 2024) pending acceptance and filing; not yet on file. Former provider: Brevo SAS (Paris, FR), used until 2026-09-26; Art. 28 DPA never executed; wind-down pending (see ROPA OI-05).

7. DPO Consultation (Art. 35(2))

Article 35(2) GDPR requires the controller to seek the advice of the Data Protection Officer when carrying out a DPIA, and to include the DPO's views in the assessment document.

❏ **DECISION NEEDED — DPO Designation and Independence:** The current privacy contact (privacy@ciata.io) is Prof. Dr. Tim Landgraf, who is also the Managing Director of CIATA GmbH. Under Art. 38(3) GDPR and EDPB Guidelines on Data Protection Officers (WP243 rev.01, §3.5), a DPO must be able to perform their duties independently and must not receive instructions regarding the exercise of those tasks. An individual who simultaneously holds

the role of Managing Director faces an inherent structural conflict of interest when assessing the lawfulness of the data processing activities that they control.

CIATA must decide between the following options before this DPIA can be considered fully Art. 35(2)-compliant:

- **Option A (recommended):** Designate an **independent DPO**. GDPR permits either an employee or an external DPO on a service contract (Art. 37(6)); an **external DPO** (specialist firm or DPO-as-a-Service) is recommended here because it removes the conflict by construction. An **internal** DPO is possible only if the person does not determine the purposes or means of processing – which excludes the Managing Director and, realistically, the ML/engineering staff who design the data pipeline (CJEU, X-FAB Dresden, C-453/21, 9 Feb 2023; EDPB WP243 rev.01). In a company of CIATA's size a conflict-free, sufficiently-expert internal candidate is hard to staff credibly, so an external DPO is the pragmatic route.
- **Option B (verify carefully – likely unavailable):** Treat DPO designation as voluntary. This is only defensible if **none** of the mandatory triggers apply – and these include the **German § 38(1) BDSG** overlay, not just Art. 37 GDPR. Under § 38(1) BDSG a DPO is mandatory **regardless of headcount** where the processing is subject to a mandatory DPIA under Art. 35 (see §1.1), or where ≥ 20 persons are constantly engaged in automated processing, or where data is processed commercially for transfer or for market/opinion research. **Because §1.1 assesses this processing as requiring a DPIA, Option B is likely foreclosed** unless counsel concludes the DPIA is not in fact mandatory. Any reliance on Option B must be confirmed in a written legal opinion that addresses § 38(1) BDSG expressly.

Until this decision is taken and documented, the DPO consultation step of Art. 35(2) cannot be verified as independent.

DPO consultation record: [To be completed: DPO name or firm; date of consultation; summary of advice or reference to separate advice document]

8. Residual Risk Rating and Art. 36 Prior Consultation

8.1 Residual Risk Summary

Risk	Post-mitigation Rating
R1 – Incidental structured PII in stored job state and LLM submission	Low-Medium
R2 – Unauthorised research report access	Low
R3 – Server-side data breach	Low – at-rest encryption not enabled (accepted gap; compensating controls incl. encrypted EU backups – A1)
R4 – Account-data breach	Low

Risk	Post-mitigation Rating
R5 — Art. 9 data in research query	Low-Medium
R6 — EEA-to-US backend transfer	Low (conditional on SEARCH_BACKEND=staan in prod — A2)
R7 — Staan full_content relay	Low
R8 — LLM hallucination	Low
R9 — Cross-border LLM transfer	N/A
R10 — Research content transmitted by report-ready email	Low

Aggregate residual risk: Low to Low-Medium. No risk registers at High after mitigation.

8.2 Art. 36 Prior Consultation

Under Art. 36(1) GDPR, the controller must consult the supervisory authority (Berliner Beauftragte für Datenschutz und Informationsfreiheit, BlnBDI) prior to processing only where the DPIA indicates that the processing would result in a **high residual risk** in the absence of additional measures (EDPB WP248 rev.01, §IV). The threshold for prior consultation is the inability to mitigate high risk by measures taken by the controller alone.

As the aggregate residual risk of this processing does not reach the threshold of "high risk" after the controls documented in Sections 2 and 4 are applied, **prior consultation under Art. 36 GDPR is not required at this time.**

This conclusion is conditional on resolution of the following open items:

1. Confirmation that the production deployment pins SEARCH_BACKEND=staan (R6 / action A2).
2. Confirmation of the independent-controller classification of Billie and Qonto, and their Art. 13/14 disclosure (Section 6 / actions A3, A4).

Regarding at-rest encryption (R3): it was confirmed on 2026-07-15 that the production disks are **not** encrypted at rest. R3 is nonetheless rated Low on the basis of the documented compensating controls (encrypted EU-geo-redundant backups Frankfurt → Helsinki, per-user access isolation, and 30-day retention); enabling volume-level at-rest encryption before the hot-production cutover remains a recommended hardening step (action A1) rather than a condition for the Low rating.

If items 1–2 cannot be confirmed or cannot be remediated, the risk rating for R6 may need to be revised upward, and prior consultation with the BlnBDI should be reconsidered.

9. Conclusion

This DPIA concludes that the processing activities of CIATA DeepResearch are **necessary and proportionate** to the legitimate purposes pursued, subject to the open action items identified below. All identified risks have been mitigated to an acceptable level (Low to Low-Medium)

through the technical and organisational measures documented in this assessment. **No Art. 36 prior consultation is required**, subject to resolution of the flagged items.

NOTE: External legal review of this DPIA has NOT yet been conducted. This document should be reviewed by an independent legal practitioner or qualified DPO before being relied upon for regulatory purposes.

Open Action Items

#	Item	Owner	Priority
A1	At-rest disk encryption confirmed not enabled (2026-07-15); accepted gap with compensating controls (encrypted EU backups, access isolation, 30-day retention). Decide whether to enable volume-level (LUKS/dm-crypt) encryption before the hot-production cutover	Engineering	High
A2	Confirm SEARCH_BACKEND=staan is pinned in production .env or config.yaml; document	Engineering	High
A3	Confirm Billie GmbH independent-controller classification with counsel; Art. 13/14 disclosure completed (privacy policy §6)	Legal / MD	Medium
A4	Confirm Qonto independent-controller classification with counsel; Art. 13/14 disclosure completed (privacy policy §6)	Legal / MD	Medium
A5	Resolve DPO designation and independence question (Section 7); document decision	MD	High
A6	Commission external legal review of this DPIA	MD / Legal counsel	Medium
A7	Confirm minimum statutory retention period for withdrawal_consent records with counsel; §§195/199 BGB limitation periods likely govern; records must NOT be deleted on account erasure or 30-day job TTL cleanup until minimum period confirmed	Legal / MD	High
A8	Confirm §147 AO interaction with company profile (company_profiles table) account-erasure obligation with counsel; fields used to generate invoices may be subject to 10-year §147 AO retention even after account deletion — erasure of those fields within the 10-year period may require redaction rather than deletion	Legal / MD	High
A9	Before enabling OpenObserve telemetry (src/perf.py, currently inert): (a) decide self-hosted vs. hosted — the choice determines whether it is a third-party sub-processor requiring an Art. 28 DPA and a transfer assessment; (b) obtain counsel's view on whether job_id-linked metric records are pseudonymous personal data; (c) decide whether to scrub or drop the raw exception error field, which cannot be categorically guaranteed free of incidental content; (d) add the disclosure to / privacy.html; (e) add entries to ROPA.md and DPA_REGISTER.md. See ZDR_STATEMENT.md §5.9	Engineering / Legal	High

Approved by: Prof. Dr. Tim Landgraf, Managing Director, CIATA GmbH **Date:** 2026-08-11