

Technische und organisatorische Maßnahmen (TOM) — Technical and Organisational Measures

CIATA DeepResearch — Annex to Data Processing Agreement (Art. 32 GDPR)

Controller / Auftraggeber	As identified in the main DPA
Processor / Auftragnehmer	CIATA GmbH, Arnimallee 7, 14195 Berlin · HRB 266084 B, Amtsgericht Charlottenburg
Product / Produkt	CIATA DeepResearch (research.ciata.io)
Document version	1.6
Date	2026-09-27
Status	DRAFT — not yet reviewed by external counsel
Privacy contact	privacy@ciata.io (no DPO designated — see §6)

How to read this document. This annex lists technical and organisational measures that are **in place today** against the codebase on main as at 2026-09-26.

Revision 1.6 (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 (§5.1) updated: Consistency Checker added to the qwen3.8-27b agent list.

Revision 1.5 (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 (gpt-oss-120b, qwen3.8-27b, gemma4-chat-26b). Failed repairs now remove leaked characters in code. Sub-processor table (§5.1) and OI-3 updated: model list for AKI.io language guard changed; "four" changed to "three". Same data categories, same recipient (AKI.io GmbH, Frankfurt).

Revision 1.4 (2026-09-26): Transactional email moves from Brevo SAS to Scaleway SAS (issues #152, #159). Brevo is noted as former processor, with its wind-down as an open item. Sub-processor table (§5.1), TLS outbound call list (§1.5), EU-only claim (§2.1), DPA status (§5.2), and OI-2 updated accordingly. Report-ready mail (issue #159) sends report title, summary, link, and deletion date via the email processor; it is on by default and can be

disabled in account settings. The production server uses a sending-only Scaleway key (TransactionalEmailEmailApiCreate). Scaleway TEM delivers links without tracking redirects.

Revision 1.3 (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).

Revision 1.2 (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.

Revision 1.1 (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. Each technical measure cites the implementing source file so that a customer's auditor can verify the claim. Where a measure is planned but not yet implemented, or where a gap exists, the status is stated explicitly. Overstated TOM annexes are contractual misrepresentations; this document does not pad.

1. Vertraulichkeit / Confidentiality

1.1 Zutrittskontrolle — Physical Access Control

CIATA DeepResearch runs exclusively on **Hetzner Cloud** (Germany). CIATA operates no physical hardware of its own. Physical datacenter access controls are governed by Hetzner's certified security programme.

Measure	Status	Notes
ISO 27001 certified datacenter	Confirmed	Hetzner Online GmbH, Germany
24/7 on-site security, CCTV, access-card zones	Confirmed	Hetzner datacenter standard
No CIATA staff physical access to servers	Confirmed	Shared virtualised cloud infrastructure

Note. Disk encryption at rest is **not enabled** on the production host. See §1.5 for the full disclosure and compensating controls.

1.2 Zugangskontrolle — System Access Control (Authentication)

All API access requires authentication. Three authentication methods are implemented and checked in order on every request (`src/api/auth.py : get_user_id()`).

Measure	Detail	Implementing file
Password hashing (bcrypt)	All user passwords stored as bcrypt hashes (adaptive work factor). Plain-text password is never written to disk or logged.	<code>src/api/auth_users.py — _bcrypt.hashpw()</code> , <code>_bcrypt.checkpw()</code>
JWT access tokens (HS256)	Short-lived (configurable; default 60 minutes). Payload carries sub (<code>user_id</code>), admin flag, and tv (token-version generation counter).	<code>src/api/auth_users.py : create_access_token()</code>
JWT refresh tokens	Longer-lived (configurable; default 7 days). Carry the same tv claim.	<code>src/api/auth_users.py : create_refresh_token()</code>
Server-side token revocation (token_version)	Every JWT-authenticated request validates the token's tv claim against a <code>token_version</code> counter stored in the user record (database schema v12). Password change (PATCH <code>/v1/auth/me/password</code>) and password reset (POST <code>/v1/auth/reset-password</code>) both call <code>db.bump_token_version()</code> , which invalidates all outstanding access and refresh tokens for that user immediately — independently of their remaining TTL. A mismatch returns HTTP 401. This closes the window in which a stolen token would remain valid after the user takes remedial action.	<code>src/api/auth_users.py : token_version_matches()</code> , lines 585 and 675; <code>src/api/auth.py : _user_id_from_jwt()</code>
Per-user API keys	Keys carry the prefix <code>dr_</code> and are bcrypt-hashed in the database. Only the raw key is shown once at creation; subsequent requests validate against the stored hash.	<code>src/api/auth_users.py : create_api_key()</code> , <code>src/api/auth.py</code>
Legacy shared API key (backward compat)	Retained for backward compatibility only; maps to the sentinel <code>user_id = "__legacy_api_key__"</code> and does not grant access to any user's research data. No new deployments should rely on this mechanism.	<code>src/api/auth.py : LEGACY_USER_ID</code>
Email verification gate	Accounts are inactive until the owner clicks a verification link (72-hour TTL). Login returns HTTP 403 for unverified accounts. Unverified accounts are auto-deleted by a background cleanup after 72 hours.	<code>src/api/auth_users.py : verify_email()</code> , <code>login()</code>
Admin role	Separate boolean flag in the user record; required for administrative endpoints. The admin role does not confer ownership of other users' research data (see §1.3).	<code>src/api/app.py : _require_admin()</code>
Rate limiting on authentication endpoints	IP-based sliding-window rate limiter (10 req/60s default, configurable). Applied to <code>/register</code> , <code>/login</code> , <code>/forgot-password</code> , and <code>/reset-password</code> .	<code>src/api/auth.py : RateLimiter</code> , <code>rate_limit_by_ip()</code>

Measure	Detail	Implementing file
JWT secret validation at startup	Application refuses to start if JWT_SECRET is unset or equals the insecure default.	src/api/auth_users.py : validate_jwt_secret()
Account-status re-check on every request	Every authenticated request verifies that user.is_active is still True. Deleted or deactivated accounts are immediately rejected (HTTP 401) regardless of token validity.	src/api/auth.py : _user_id_from_jwt()

1.3 Zugriffskontrolle — Data Access Control (Authorisation)

Measure	Detail	Implementing file
Per-endpoint ownership verification	Every endpoint that returns or deletes research data calls _check_job_ownership(). This raises HTTP 404 if the authenticated user_id does not match the job's recorded owner.	src/api/app.py : _check_job_ownership() (line 279)
No administrative override on user data	_check_job_ownership() contains no bypass for the admin role. CIATA staff cannot read, download, or delete another user's research run through any API endpoint. This includes the admin report inspector (GET /v1/admin/reports/{job_id}/sources/{i}/text), which requires both the admin role and job ownership.	src/api/app.py : _check_job_ownership() and _require_admin() — called independently; ownership is not waived by admin status
Research submission rate limiting	Per-user limits on concurrent jobs and hourly job submissions, checked before a new job is created.	src/api/rate_limit.py : check_research_rate_limits()
API-level rate limiting	Sliding-window per-API-key rate limiter (30 req/60s default). Returns HTTP 429 with a Retry-After header.	src/api/auth.py : RateLimiter
GDPR Art. 17 deletion endpoints	DELETE /v1/research/{job_id} — immediate deletion of a single run, with ownership verification. DELETE /v1/auth/me — account soft-delete requiring password confirmation; clears PII and deactivates the account.	src/api/app.py, src/api/auth_users.py : delete_account()
In-flight erasure support	The research worker re-checks user.is_active before a job starts and at each research iteration. A queued or in-progress job is aborted — not stored — when the owning account is deleted.	src/api/worker.py, src/engine.py (account-status re-check per iteration)

1.4 Trennungskontrolle — Separation of Data

Measure	Detail	Implementing file
Per-user data isolation (filesystem)	Each research run is stored in its own directory <code>./runs/{job_id}/</code> . The <code>job_id</code> is a UUID assigned at submission.	<code>src/api/job_store.py</code>
Per-user data isolation (database)	All SQLite tables carry a <code>user_id</code> foreign key. No query returns data across user boundaries at the application layer.	<code>src/api/database.py</code>
Single-tenant product scope	DeepResearch is a single-deployment SaaS product. There is no multi-tenancy at the organisation level; each user account is isolated by <code>user_id</code> .	Architecture (no code ref)
Non-root container	The Docker container runs as UID 1000, limiting the blast radius of any container-escape vulnerability.	Dockerfile : USER 1000 (line 71)
Separate Docker volumes	Research data (<code>./runs</code>) and the SQLite database (<code>./data</code>) are mounted as separate named volumes, isolating them from the application layer.	<code>docker-compose.yml</code>

1.5 Verschlüsselung / Encryption — At Rest and In Transit

In Transit

Measure	Detail	Notes
TLS (HTTPS) for all client traffic	Caddy reverse proxy handles TLS termination with automatic certificate provisioning via Let's Encrypt. All communication between the user's browser or API client and the server is encrypted in transit.	Infrastructure — Caddy configuration
TLS for sub-processor calls	All outbound API calls (AKI.io LLM, Staan search, Mollie payments, Scaleway TEM) use HTTPS.	Application-level (httpx, SDK clients)
Scaleway TEM: sending-only API key, no link tracking	The production server authenticates to Scaleway TEM using an IAM permission set scoped to <code>TransactionalEmailEmailApiCreate</code> only (no read, delete, or administrative permissions). Scaleway TEM has no link-tracking feature; links in outbound mails are delivered exactly as written, without redirect through any tracking domain.	<code>src/api/email_service.py</code> (Scaleway sender)

At Rest

Measure	Status	Detail
Volume encryption at rest (Hetzner)	NOT ENABLED	The production <code>./runs</code> and <code>./data</code> Docker volumes are not encrypted at rest. No LUKS, dm-crypt, or Hetzner volume-level encryption is active on the production host (confirmed 2026-07-15). This is a documented, accepted gap.

Measure	Status	Detail
Compensating controls (in place now)	Implemented	(a) bcrypt-hashed passwords — if the database file were exfiltrated, no plaintext passwords would be exposed; (b) TLS in transit prevents interception before data reaches disk; (c) strict per-user application-level access isolation — an attacker who obtained a database file would still need a valid session to query research data via the API; (d) automated 30-day retention reduces the live data set at any point in time; (e) encrypted off-site backups held geo-redundantly within the EU (Hetzner Frankfurt → Helsinki).
Backups	Encrypted, EU geo-redundant	Hetzner Storage / backup service. Backups are encrypted and held geo-redundantly across Hetzner's Frankfurt and Helsinki regions. CIATA does not hold unencrypted copies outside Germany.
Enabling volume-level at-rest encryption	Recommended hardening — not yet implemented	Enabling LUKS/dm-crypt encryption on the production volumes before the hot-production cutover is a recommended hardening step (DPIA action A1). A customer's DPO should treat this as an open item and may require a timeline for remediation.

Directly answering the question a DPO will ask: Data at rest on the production server is **not** encrypted by a volume- or block-device-level mechanism. The database file (`data/*.db`) and research run state files (`runs/{job_id}/state.json`, `report.md`) reside on unencrypted Hetzner Cloud disk volumes. The compensating controls listed above are the operative safeguards until volume-level encryption is enabled.

1.6 Pseudonymisierung — Pseudonymisation

Measure	Detail	Implementing file
PII filter on all scraped content	A deterministic regex-based filter (<code>scrub_pii()</code>) is applied to all web page text before it is written to disk and before it is sent to any LLM. The filter runs at <code>src/engine.py:859</code> — before <code>source.page_text</code> is assigned — so neither the stored <code>state.json</code> field nor any LLM prompt ever contains raw, unfiltered page content.	<code>src/pii_filter.py : scrub_pii()</code>
PII filter coverage	The filter replaces the following with placeholder tokens: email addresses [EMAIL], phone numbers in international and DACH formats [TELEFON], German Steuerliche Identifikationsnummer (11-digit) [STEUER-ID], German Sozialversicherungsnummer [SOZIALVERS-NR], Swiss AHV/AVS numbers [AHV-NR], IBANs [IBAN], Luhn-validated credit card numbers [KREDITKARTE], IPv4 addresses [IP-ADRESSE], and German-format postal addresses [ADRESSE].	<code>src/pii_filter.py : _PATTERNS</code> (lines 101–111)
Known PII filter limitations (residual risk — stated honestly)	The regex filter does not catch: personal names mentioned in running text, non-DACH national identifiers (UK National Insurance numbers, US Social Security Numbers, French/Dutch/Polish national ID formats), or postal addresses in non-German formats. Research questions submitted by users may	DPIA §3.3, R1

Measure	Detail	Implementing file
	also contain names and other personal data that CIATA cannot screen. This residual is accepted as proportionate under Art. 6(1)(f) (Legitimate Interest) and is documented in the DPIA (§3, R1).	
Question hash / question_hash	The first 16 hexadecimal characters of a SHA-256 hash of the research question are embedded in each report's machine-readable AI-provenance comment (see §4.5). It is a one-way value that distinguishes one report from another; the question cannot be reconstructed from it. It is derived at render time and is not stored as a separate field in the job record, and is not used for deduplication.	src/report.py — computed in render_report()
Input sanitisation (watermark injection prevention)	User-supplied research questions are stripped of VSRMark variation selectors (U+FE00–FE0E, U+E0100–E01EF) and zero-width formatting characters (U+200B–200D, U+2060, U+FEFF) at API ingestion. This prevents a caller from injecting steganographic characters that would corrupt the AI-provenance watermark of the resulting report.	src/watermark.py : sanitize_input()

2. Integrität / Integrity

2.1 Weitergabekontrolle — Transfer Control (Data in Motion)

Measure	Detail	Notes
TLS on all data paths	All traffic between users and the server, and between the server and sub-processors, is encrypted in transit using TLS.	Caddy (client-facing); HTTPS via httpx/SDKs (sub-processors)
EU-only sub-processors (default configuration)	In the default production configuration (SEARCH_BACKEND=staan), all data remains within the EU/EEA: Hetzner (DE), AKI.io (Frankfurt, DE), Staan/European Perspective SAS (EU), Mollie (NL), Scaleway (FR, fr-par only), Billie (DE), Qonto (FR). No personal data leaves the EEA in the default configuration.	See §5 for sub-processor list
Non-default backends carry cross-border transfer risk	The system supports DuckDuckGo (US) and Google Custom Search (US) as alternative search backends, enabled via the SEARCH_BACKEND configuration variable. If activated in production, these constitute EEA-to-US transfers without Standard Contractual Clauses currently in place. Staan must remain the pinned production backend to preserve the EU-only claim.	src/searcher.py, ZDR_STATEMENT.md §6
No AI training on user data	The DPA with AKI.io (when executed) will prohibit use of CIATA prompts and completions for model training. Until the DPA is executed, this is an intended contractual commitment, not a confirmed one. See §5.	DPA pending — see §5

Measure	Detail	Notes
LLM input is PII-filtered	Page content sent to AKI.io models is filtered by <code>scrub_pii()</code> before the LLM call, reducing the personal data surface in transit to AKI.io infrastructure.	<code>src/pii_filter.py</code> , <code>src/engine.py</code>
Derived search queries only	Search queries sent to Staan are generated by the LLM from the user's question; the original research question is not transmitted to the search provider.	<code>src/engine.py</code> (query generation step)

2.2 Eingabekontrolle — Input Integrity and Logging

Measure	Detail	Implementing file
Request ID on every request	A UUID-based X-Request-ID header is generated (or forwarded if supplied by the caller) and attached to every request and response. The ID is available to all downstream log entries, enabling end-to-end correlation of a single request across server logs.	<code>src/api/middleware.py</code> : <code>RequestIDMiddleware</code>
Structured access logging	Every completed HTTP request is logged with method, path, HTTP status code, wall-clock duration (ms), and request ID at INFO level.	<code>src/api/middleware.py</code> : <code>AccessLogMiddleware</code>
Unhandled exception logging	Any exception that escapes the endpoint stack is caught, logged with full traceback at ERROR level (including request ID), and returns a generic HTTP 500 to the caller without leaking internal details.	<code>src/api/middleware.py</code> : <code>global_exception_handler()</code>
Input sanitisation (research questions)	Research questions are validated for length (maximum 4,000 characters per Pydantic model) and stripped of steganographic characters before reaching the engine.	<code>src/api/models.py</code> (Pydantic schema), <code>src/watermark.py</code> : <code>sanitize_input()</code>
SSRF protection on outbound fetches	Before fetching any URL, <code>is_safe_url()</code> validates that: (a) the scheme is http or https; (b) the hostname resolves; (c) no resolved IP falls within private, loopback, link-local, or reserved ranges. This prevents server-side request forgery via crafted URLs in search results.	<code>src/url_validation.py</code> : <code>is_safe_url()</code>
robots.txt compliance	The page reader checks the target site's <code>robots.txt</code> before every direct CIATA fetch. Disallowed URLs are skipped. Per-domain cache TTL: 1 hour. Graceful-fallback if <code>robots.txt</code> is unreachable. Scope note: in the production-default Staan <code>full_content</code> mode, Staan fetches most page text server-side and CIATA's direct <code>robots.txt</code> check does not run for those results; CIATA only fetches directly when Staan returns no content.	<code>src/page_reader.py</code> (lines 44–113)
Honest bot User-Agent	All direct CIATA page fetches identify themselves as a bot with the contact address: CIATA-DeepResearch/1.0 (+https://research.ciata.io; bot; contact:	<code>src/page_reader.py</code> : <code>USER_AGENT</code> (line 31)

Measure	Detail	Implementing file
	privacy@ciata.io). The User-Agent is never spoofed as a browser.	
Per-domain rate limiting on page fetches	A minimum of 1 second is enforced between consecutive direct fetches to the same domain (thread-safe).	src/page_reader.py (lines 50, 116–129 area)
Content truncation	Fetched page text is truncated to a maximum of 16,000 characters per source in production (code default: 8,000 characters, overridden in config.yaml).	src/page_reader.py (line 194–195 area)

3. Verfügbarkeit und Belastbarkeit / Availability and Resilience

Measure	Detail	Notes
Container health check	GET /v1/health runs inside the container every 30 seconds with a 5-second timeout and 3-retry policy. The endpoint verifies database connectivity and sub-processor reachability. Docker restarts the container if health checks fail repeatedly.	Dockerfile : HEALTHCHECK
Automatic container restart	restart: unless-stopped in docker-compose.yml ensures the API server restarts automatically after any unexpected crash.	docker-compose.yml
Graceful shutdown	On SIGTERM, the container waits up to 300 seconds for in-flight research jobs to complete before terminating. Research state is saved incrementally during the loop, so an interrupted job can be diagnosed from its last save.	docker-compose.yml : stop_grace_period: 300s
Encrypted off-site backups	Hetzner's volume backup service stores encrypted backups geo-redundantly across EU regions (Frankfurt primary → Helsinki secondary). CIATA does not hold unencrypted off-site copies.	Hetzner infrastructure; confirmed in LAW_FIRM_BRIEFING §3.3
Automated data retention	A background timer runs every 24 hours and deletes run directories (state.json, report.md) older than DATA_RETENTION_DAYS (default: 30 days). This limits the live data set and reduces the storage surface. The retention clock is the created_at timestamp in job_meta.json.	src/api/job_store.py : _get_run_created_at(), retention cleanup logic
Thread pool for research workers	The research engine runs in a thread pool (maximum 10 concurrent workers). Each research job is isolated in its own thread.	src/api/worker.py (max_workers=10)
Credit refund on unrecoverable errors	If the language guard cannot repair a corrupted report, the job is failed and the user's credit is refunded. This prevents charging users for unusable output.	src/engine.py (language guard fail path)

4. Verfahren zur regelmäßigen Überprüfung, Bewertung und Evaluierung / Regular Review, Assessment, and Evaluation

4.1 Testing

Measure	Detail	Notes
Automated test suite	2,231 tests, all passing on main as at 2026-08-11. Tests cover unit, integration, and evaluation scenarios.	LAW_FIRM_BRIEFING.md v1.2 header
Per-deploy test gate	Tests are run before deployment; the production container is not rebuilt from a branch where tests fail.	Engineering practice

4.2 Code Review and Deployment

Measure	Detail	Notes
Git-based deployment	All changes to the system are tracked in version control. The deployed container is built from a tagged commit. Run provenance metadata (code_revision — short git SHA, agent_models, config_digest) is embedded in every research run's state.json so that any run can be traced to the exact code and configuration that produced it.	docker-compose.yml, src/engine.py (provenance block)
Staged deployment	The scripts/deploy.sh deployment script captures the git revision at deploy time (GIT_REVISION) and stamps it into the container via the VERSION file.	Dockerfile : ARG GIT_REVISION, scripts/deploy.sh

4.3 Compliance Review

Measure	Detail	Notes
DPIA	A Data Protection Impact Assessment has been conducted and is maintained at docs/compliance/DPIA.md. Annual review cycle (next: 2027-06-18).	DPIA v1.2, 2026-08-11
LIA	A Legitimate Interest Assessment for web scraping is maintained at docs/compliance/LIA.md.	LIA v1.2, 2026-08-11
ZDR Statement	A Zero Data Retention Statement documenting sub-processor retention positions is maintained at docs/compliance/ZDR_STATEMENT.md.	ZDR v2.1, 2026-08-11
External legal review	This TOM annex, the DPIA, and the DPA have not yet undergone external legal review . Review by qualified data-protection counsel is a pre-launch open item.	Open item

4.4 Incident Response

Measure	Detail	Notes
Incident response	A written breach-response procedure exists: docs/compliance/INCIDENT_RESPONSE.md (v1.0, 2026-08-11). It defines the 72-hour	Drafted — review and

Measure	Detail	Notes
procedure (Art. 33/34 GDPR)	clock, a three-outcome severity triage, containment steps specific to this stack, notification templates for the supervisory authority and for data subjects, and the Art. 33(5) breach register. It is drafted but not yet externally reviewed or rehearsed , and several roles currently rest on a single person.	rehearsal outstanding, see §6

4.5 EU AI Act Transparency Measures (Art. 50(2))

Requirement	Measure	Implementing file
Machine-detectable AI watermark	Every research report carries a VSRMark steganographic watermark encoded as Unicode variation selectors. The payload is JSON: {"producer": "CIATA Deep Research", "version": "1.0", "ai_generated": true, "job_id": "...", "timestamp": "..."}.	src/watermark.py : watermark_report()
Machine-readable HTML provenance comment	An HTML comment block is embedded in every report: <!-- AI-PROVENANCE generator: CIATA DeepResearch/1.0 ... eu_ai_act: Article 50(2) transparency disclosure -->.	src/report.py (lines 150–164 area)
Visible AI disclosure	Every report contains the plaintext notice: "This report was generated by AI (CIATA DeepResearch). It may contain errors or omissions. Verify critical facts independently."	src/report.py (lines 194–197 area)
PDF AI disclosure banner	Branded PDF exports include an AI disclosure banner in the header and a data-processing footer on every page.	src/pdf_export.py (lines 171–204 area)
C2PA content provenance (PLANNED — NOT IMPLEMENTED)	C2PA content credentials on PDF output are planned for a post-launch release. This measure must not be treated as present.	Planned; DPIA §4.1

5. Auftragskontrolle / Processor Control — Sub-processors

5.1 Sub-processor List

CIATA engages the following sub-processors for the processing covered by the DPA to which this annex is attached. All named sub-processors operate exclusively within the EU/EEA in the **default production configuration**.

Sub-processor	Role	Location	Data transmitted	Art. 28 DPA status
AKI.io GmbH	LLM inference — three models: gpt-oss-120b (reasoning, relevance gate), qwen3.8-27b (synthesis, table repair, paragraph check, consistency check, patch), gemma4-chat-26b (extraction and language guard, reasoning disabled). qwen3.6-chat-35b was used for the language guard until 2026-09-27.	Frankfurt, Germany DE	Research questions; PII-filtered page excerpts (max 16,000 chars/source); LLM prompts and completions	⚠ Required — not yet executed / not on file
European Perspective SAS (Staan)	Web search (Ecosia + Qwant index); in full_content mode (production default) Staas server-side fetches result pages and returns full text, which CIATA then PII-filters	EU EU	Derived search queries (not original question); in full_content mode: full page text (markdown) of search results	⚠ Required — not yet executed / not on file
Hetzner Online GmbH	Cloud infrastructure hosting; encrypted EU geo-redundant backups	Germany DE	All application data (at infrastructure/hypervisor level). Disk not encrypted at rest — see §1.5.	⚠ Required — not yet executed / not on file
Mollie B.V.	Card payment processing	Amsterdam, Netherlands NL	Payment amount, currency; CIATA-internal user_id; billing address if provided	⚠ Required — not yet executed / not on file
Scaleway SAS	Transactional email (account verification, password reset, purchase confirmation, report-ready notification). Active since 2026-09-26 (issue #152). Former processor Brevo SAS remains in code as rollback only; after the switch no personal data is sent to Brevo.	fr-par, Paris, France FR	Email address, name; purchase details; verification / reset tokens; report-ready mail: additionally report title, a short summary, link, and report deletion date	⚠ Required — not yet executed / not on file (Scaleway DPA "Version of June 1st, 2024" pending acceptance)

Independent controllers (no Art. 28 DPA required — Art. 13/14 disclosure provided)

Party	Role	Location	Notes
Billie GmbH	B2B invoice / Net-30 BNPL (embedded in Mollie checkout; activated only when buyer	Berlin, Germany DE	Receives billing name/address, VAT-IdNr., Handelsregister number for credit

Party	Role	Location	Notes
	selects Billie as payment method)		underwriting. Independent controller classification to be confirmed with counsel.
Qonto SAS	Invoice generation — called for every completed purchase	Paris, France 	Receives buyer name, email, billing address, VAT number. Independent controller classification to be confirmed with counsel.

Not a sub-processor

Plausible Analytics — self-hosted by CIATA at trc.ciata.io (Hetzner Cloud, Germany). Cookie-free; no personal data is collected or processed. CIATA operates the server; no data is transmitted to Plausible's own infrastructure. This is not a third-party sub-processor.

5.2 Sub-processor Governance

Measure	Status
Prior written authorisation from controller before engaging new sub-processors	Committed — to be included in the DPA
Sub-processor DPA chain (Art. 28(4))	Intended for all five processor-role sub-processors listed above
Notification of sub-processor changes	Committed — notification period and mechanism to be agreed in the DPA
DPAs currently executed	None. As of 2026-09-26, no sub-processor DPA (AKI.io, Staan, Hetzner, Mollie, Scaleway) has been formally executed or is on file. This is a pre-launch open item. Staan's standard DPA is on file since 2026-10-08 (it keeps queries for 25 months). Zero-retention and no-training commitments from AKI.io and Staan are intended contractual terms, not yet confirmed by a signed agreement.

6. Offene Punkte / Open Items

The following measures are recommended but not yet implemented, or require confirmation before they can be stated as present. A customer's DPO should treat this as the list of items requiring a resolution timeline before the DPA takes effect.

#	Item	Priority	Status
OI-1	Volume-level at-rest encryption (LUKS/dm-crypt on the production Hetzner volumes). Enables removal of the at-rest encryption gap disclosure in §1.5.	High	Not implemented (DPIA action A1)
OI-2	Execute all five sub-processor DPAs (AKI.io, Staan, Hetzner, Mollie, Scaleway). Until executed, zero-retention and no-training commitments are not contractually confirmed.	High	None executed

#	Item	Priority	Status
	Scaleway DPA is "Version of June 1st, 2024"; acceptance pending. Former processor Brevo's DPA was never executed; wind-down and log-deletion request remain open.		
OI-3	Confirm AKI.io DPA covers all three model IDs currently in use (gpt-oss-120b, qwen3.8-27b, gemma4-chat-26b) and past processing on qwen3.6-chat-35b (in use until 2026-09-27).	High	Pending DPA execution
OI-4	Review and rehearse the incident-response procedure (docs/compliance/INCIDENT_RESPONSE.md, v1.0). The document is drafted and covers Art. 33 and Art. 34 notification. Outstanding: external legal review, a first tabletop rehearsal, creation of the breach register file, and appointment of a deputy — every incident role currently sits with the Managing Director alone.	High	Drafted — not reviewed or rehearsed
OI-5	DPO designation. No Data Protection Officer has been designated. Under § 38(1) sentence 2 BDSG, a DPO is mandatory regardless of headcount where processing is subject to a mandatory DPIA (Art. 35 GDPR) — which this processing is assessed to be. The Managing Director (privacy@ciata.io) currently serves as the privacy contact, but this person simultaneously controls the processing and therefore cannot perform the DPO function independently (Art. 38(3) GDPR; EDPB WP243 rev.01). An external DPO is recommended.	High	Not resolved (DPIA §7)
OI-6	External legal review of this TOM annex and the DPIA. Neither document has been reviewed by qualified external data-protection counsel.	Medium	Not completed
OI-7	Confirm Billie and Qonto independent-controller classification with counsel. If either is found to be a processor, an Art. 28 DPA must be obtained.	Medium	Pending legal confirmation
OI-8	Confirm minimum retention period for withdrawal-consent records (withdrawal_consents table) under §§ 195/199 BGB limitation periods. These records must not be deleted with the 30-day job TTL or on account erasure until the minimum period is confirmed.	High	Pending legal confirmation (DPIA action A7)
OI-9	Confirm §147 AO interaction with company billing profile (company_profiles table) — fields used to generate invoices may be subject to a 10-year §147 AO retention obligation even after account deletion.	High	Pending legal confirmation (DPIA action A8)
OI-10	Confirm Art. 50(2) EU AI Act compliance of the VSRMark watermark. VSRMark is a proprietary Unicode steganographic scheme. Whether a non-standardised watermark discharges the Art. 50(2) machine-detectable marking obligation (which has applied since 2025-08-01) requires a legal opinion. C2PA content provenance is planned post-launch and would address this more robustly.	Medium	Open legal question (LAW_FIRM_BRIEFING §6.2)

#	Item	Priority	Status
OI-11	Confirm production SEARCH_BACKEND=staan is pinned in the production .env or config.yaml. If a US-based backend (DuckDuckGo, Google CSE) is active, the EU-only data-transfer claim in §2.1 does not hold and Standard Contractual Clauses must be in place before that backend is used.	High	Requires deployment verification

End of TOM annex. This document reflects the system on main as at 2026-08-11. It should be updated whenever a measure is added, modified, or an open item is resolved.

Approved by: Prof. Dr. Tim Landgraf, Managing Director, CIATA GmbH

Date: 2026-09-13

Status: Draft — external legal review pending