Legislator Privacy Policy

Last updated: 28 October 2026

1. Scope and responsible company

This Privacy Policy explains how **Legislator AI AG**, Kirchstrasse 6, 8807 Freienbach SZ, Switzerland, company identification number (UID) CHE-476.609.141 ("Legislator", "we", "us", or "our"), processes personal data in connection with legislator.ch, the Legislator legal AI platform and our related services (together, the "Services"). You can contact us regarding data protection and professional secrecy matters at **legal@legislator.ch**.

This Policy covers website visitors, prospective customers, individual users, representatives of customer organizations and other individuals whose personal data is processed through the Services. A "Customer" is the person or organization (such as a law firm, notary office, corporate legal department, or other entity) that obtains or uses the Services. "Customer Content" means all material submitted to, stored in, or generated through the platform on the Customer's behalf, as described in Section 3.

We process personal data in accordance with the Swiss Federal Act on Data Protection of 25 September 2020 (FADP; SR 235.1) and the Ordinance on Data Protection of 31 August 2022 (ODP; SR 235.11), as well as the EU General Data Protection Regulation (Regulation (EU) 2016/679, GDPR) where it applies to the relevant processing. This Policy fulfills our statutory transparency duties (Art. 19 FADP; Arts. 13 and 14 GDPR). It does not, by itself, replace the applicable Data Processing Agreement and Professional Secrecy Addendum ("DPA").

2. Our roles and the customer relationship

2.1. Processing for our Own purposes (Legislator as Controller)

We act as a data controller (Art. 5 lit. j FADP; Art. 4 GDPR, where applicable) for the personal data we use to administer user accounts and subscriptions, manage our business relationships and billing, operate and secure our website and Services, handle administrative support, comply with applicable legal and accounting obligations and protect our rights.

2.2. Processing on behalf of customers (Legislator as processor and professional auxiliary)

We act as a data processor (Art. 5 and Art. 9 FADP; Art. 28 GDPR, where applicable) and, in relation to legal professionals bound by professional secrecy, as an auxiliary subject to professional secrecy (Hilfsperson within the meaning of Art. 321 para. 1 of the Swiss Criminal Code (SCC) and Art. 13 para. 2 of the Swiss Federal Act on the Free Movement of Lawyers BGFA/LLCA) when handling personal data in Customer Content on a Customer's documented instructions. The Customer remains the data controller and determines the purposes and essential means of that processing. We determine and implement the technical and organizational arrangements needed to provide the agreed Services strictly within those instructions and the DPA. Where a Customer itself acts as a processor for another controller, we act as its subprocessor as appropriate.

The Customer's instructions include its agreed service configuration, the tasks, prompts, and files submitted by authorized users and its retention, sharing and deletion settings. We do not acquire any independent right to reuse Customer Content merely because it is submitted to the Services. Content made available for technical support remains subject to the same strict restrictions; administrative correspondence about a support request may be processed separately as controller data.

The DPA sets out the binding contractual arrangements for documented instructions, professional secrecy (Art. 321 SCC / Art. 13 BGFA), data security, subprocessor authorization and oversight, assistance with data subject rights, security incident notification and the return or deletion of Customer Content. **A signed copy of the DPA and the Professional Secrecy Addendum can be requested at legal@legislator.ch**. This Policy does not reduce our

duties or any more protective commitments agreed with a Customer in a signed DPA, which takes precedence over this Policy.

3. Categories of personal data processed and their sources

3.1. Account and administrative data (Controller)

Depending on your interaction with us, we process the following categories of personal data:

Identity and contact details: Name, surname, (professional) email address, organization or law firm name, professional role and contact information you provide.
Account and subscription information: User and organization identifiers, access permissions, authentication and session information, selected subscription plan, subscription status and account settings (including diagnostic review preferences).
Billing and payment information: Billing details, invoices, transaction references, payment status and subscription payment history. Stripe processes payment card credentials directly through its payment services.
Communications: Enquiries, demo requests, administrative support correspondence, voluntary feedback and communication preferences.
Technical and usage metadata: IP addresses, browser and device information, timestamps, login, security and access events, error codes and measurements of service consumption (such as token counts and computation costs). Metadata associated with an identifiable user or account is treated as administrative personal data even when it contains no document text or prompts.

We obtain this information from you, your organization or its authorized users, through the technical operation of the Services and from providers involved in authentication, communications or payments to the extent relevant to their service.

3.2. Customer content (Processor)

Customer Content includes all data submitted to the Services for processing across all platform modes (including agentic modes such as Instant, Standard, Deep Research, Deep Research Fast, Deep Research Ultra, Memo, search modes such as Legislation and Court Decisions, and the AI Detector feature), namely: prompts, questions, instructions, uploaded documents (PDF, Word, images, scanned files, audio files), extracted text (including via OCR and audio transcription), case descriptions, notes, templates, private knowledge sources, internal messages, associated file metadata, generated answers, drafts and other outputs. It also includes saved tasks, conversation histories, revisions and vector representations (embeddings) derived from that material to provide search and retrieval functions within the Customer's isolated environment.

Customer Content may identify individuals who do not use Legislator directly, including clients, employees, counterparties, witnesses, experts and officials mentioned in legal files. It may contain particularly sensitive personal data under Art. 5 lit. c FADP (e.g., health data, data relating to criminal, tax or administrative proceedings), special categories of personal data under Art. 9 GDPR or data relating to criminal convictions and offences under Art. 10 GDPR. The Customer determines what it submits and is responsible for meeting the legal and professional conditions applicable to those data.

3.3. Public legal sources

Our legal reference database contains publicly available Swiss legal materials (federal and cantonal laws, published court decisions and legal literature) and associated citation information. Personal data appearing in those materials, where present, originates from the relevant published sources. A Customer's private documents and private knowledge sources remain strictly isolated as Customer Content; they are never added to the public reference database or made available to other Customers.

4. Purposes and legal grounds for processing as controller

We use account and administrative data to establish and manage accounts, provide platform access, administer subscriptions, calculate usage charges, process payments and respond to enquiries (Arts. 6 and 31 para. 2 lit. a

FADP; where the GDPR applies, Art. 6 GDPR for the performance of a contract or pre-contractual steps, and Art. 6 GDPR for our legitimate interest in providing and administering the organization's Services).

We process technical and usage metadata to maintain infrastructure availability, diagnose faults, secure accounts and systems, prevent misuse or unauthorized access and monitor service performance (Arts. 7 and 8 FADP; Art. 6 GDPR). These activities never authorize the reuse of Customer Content for model training or development.

We handle business correspondence and send essential account, security, billing and contractual notice messages as a necessary part of providing the Services. Where we send optional product updates, we rely on consent where required or, where permitted by applicable law (Art. 6 GDPR), on our legitimate interest in informing business contacts about relevant Legislator services. You may unsubscribe or object at any time using the message instructions or by contacting **legal@legislator.ch**.

We process data required for accounting, compliance (in particular the 10-year retention duty under Art. 958f of the Swiss Code of Obligations (CO)), and the establishment, exercise or defence of legal claims (Art. 31 para. 1 FADP; Art. 6 GDPR). Finally, we maintain and index publicly available legal materials based on our legitimate interest in providing reliable legal research functions (Art. 31 para. 1 FADP; Art. 6 GDPR), balanced against the rights of the individuals concerned.

5. How we process customer content as processor

We process Customer Content solely to carry out the tasks and documented instructions requested by the Customer. This includes:

Receiving and preparing uploaded files;
Extracting and organizing relevant text and information (which executes locally on our Swiss servers for standard digital documents, and via TLS-encrypted API calls to Google Cloud Gemini models in the EU under Zero Data Retention terms for optical character recognition (OCR) on scanned images and for audio transcription);
Executing semantic search, embedding and reranking functions (which run entirely on our proprietary Swiss infrastructure);
Carrying out generative AI inference and producing the requested outputs;
Storing the tasks, drafts and histories that the Customer chooses to retain on our Swiss servers, and executing sharing or deletion instructions.

We limit the material transmitted to an external AI inference provider to what is necessary for the requested task. We never use Customer Content for advertising, independent profiling of individuals (Art. 5 lit. f and g FADP) or building datasets for other Customers.

To ensure platform reliability and remediate technical execution errors (specifically when a task fails to execute due to an application crash or attachment-processing error or when a user flags an agent mistake in a follow-up message), we operate an automated internal diagnostic workflow (AI Review): flagged tasks are first automatically anonymized by an AI model (such as Gemini under EU Zero Data Retention terms), original file attachments are stripped and user identifiers are removed so the record cannot be linked to its owner; the anonymized trace is encrypted with a dedicated key and analyzed by a diagnostic AI agent (such as Gemini or Claude) to generate a brief diagnostic message reviewed exclusively by our Head of IT for code remediation. Anonymized diagnostic records are purged immediately once code evaluation confirms that the issue is resolved and in any event no later than 7 days after creation. When the “Delete after 24 hours” setting is enabled, task data is never accessed, reviewed or used for diagnostic purposes, including to investigate or resolve technical execution errors and no diagnostic copies are retained after the task is deleted, whether manually or automatically.

6. No use of customer content for model training

We do not use Customer Content - including prompts, uploaded files, OCR/audio-extracted text, retained histories and generated outputs - to train, fine-tune or improve our own or any third-party AI models. Our binding agreements with AI inference providers (such as Google Cloud EMEA Limited and TensorX Ltd.) strictly prohibit such use of

Customer Content processed on our behalf and we never authorize any provider to opt Customer Content into model training.

7. Infrastructure in Switzerland and AI Inference in the European Union

7.1. Our physical infrastructure in Switzerland and ingress gateway

We operate the core platform, all persistent platform storage (such as the account database, stored Customer Content, uploaded files, templates, private knowledge sources, histories and backups) and our Swiss legal reference database on **our own physical servers in Switzerland**, which are owned and exclusively controlled by Legislator AI AG. Search, embedding and reranking models run entirely on our Swiss hardware; we do not use third-party cloud storage hosts or external vector database services for those components. At the network ingress stage, encrypted client traffic (TLS) passes through a web reverse-proxy gateway operated by Amazon Web Services EMEA SARL located in Switzerland (Zurich Region) directly to our on-premises Swiss servers, with zero storage of Customer Content at the gateway.

For legal clarity under the Swiss Bar Association guidelines (SAV/FSA Guidelines on Artificial Intelligence, Section 2.1 para. 2 no. 2), although Legislator AI AG owns the underlying physical servers in Switzerland, a law firm's use of the Services constitutes external IT outsourcing rather than an internal law-firm on-premises installation, because data leaves the law firm's internal network.

7.2. Transient AI Inference in the European Union under Zero Data Retention (ZDR)

For generative AI tasks (as well as OCR and audio transcription), we engage **Google Cloud EMEA Limited** for the Gemini and Anthropic Claude model families in our approved enterprise configuration and **TensorX Ltd.** for independently hosted open-weights models, including DeepSeek. All AI inference endpoints used for Customer Content are located exclusively in the **European Union**. Retries during outages or quota events are strictly routed to the same model provider within the same EU region; our application layer blocks any fallback or routing outside the approved EU configuration (fail-closed).

Our inference configuration is governed by enterprise **Zero Data Retention (ZDR)** commitments and approved Google Cloud abuse-monitoring prompt-logging exemptions (confirmed on 1st September 2026 across all Legislator AI AG Google Cloud projects):

All data travels over TLS-encrypted connections.
Prompts, document content and generated outputs are processed solely transiently in volatile memory (RAM) to compute the response and are immediately purged from memory upon request completion, with **0 days of persistent content retention** on provider disks.
Content logging, abuse-monitoring prompt logging and provider-side content caching incompatible with ZDR are disabled.
TensorX Ltd. hosts an independent copy of the relevant open-weights model (DeepSeek) on dedicated GPU infrastructure in Ireland and Finland; selecting that mode never transmits Customer Content to the model's original developers. It offers **Zero Data Retention (ZDR)** commitments for inference and does not log prompts.

8. External web research (Google Grounding)

External web research is separate from searches of Legislator's Swiss legal reference database on our Swiss infrastructure. When a Customer uses the Deep Research function, the deployed **Google Cloud Web Grounding for Enterprise** service receives the search queries and any contextual information included in that request. Web Grounding for Enterprise provides contractual Zero Data Retention guarantees (no storage of customer data and no use for model training, unlike standard web search services).

9. Recipients and subprocessors

Access to personal data is restricted to those who need it for the purposes described in this Policy. Recipients include authorized Legislator personnel, the service providers identified in **Annex 1** and, where the Customer uses sharing

features, the organization users authorized by the Customer. Providers processing Customer Content on our behalf (such as Google Cloud EMEA Limited, TensorX Ltd. and, solely for encrypted network ingress in Switzerland, Amazon Web Services EMEA SARL) act as subprocessors under binding data processing and confidentiality agreements (Art. 9 FADP, Art. 7 ODP, Art. 28 GDPR).

By accepting the Terms of Service, the Customer grants Legislator AI AG general authorisation to engage subprocessors, including to add or replace subprocessors, to process personal data on the Customer's behalf. Each subprocessor must be bound by the same substantive data protection and confidentiality obligations as those applicable to Legislator, insofar as relevant to the processing entrusted to that subprocessor. Legislator remains liable to the Customer for its subprocessors' compliance with those obligations.

For administrative auxiliary services that never receive Customer Content, prompts or uploaded legal files:

Mailgun and Resend (Plus Five Five, Inc.) provide administrative email delivery (login codes and notifications). Task notification and sharing emails contain secure access links only (link-only emails) and never contain prompts, uploaded documents, or generated legal text.
Stripe handles subscription payments and billing data (acting as a processor for commissioned payment transactions and as an independent controller for its own fraud-prevention, regulatory, and compliance obligations).

10. International data transfers and safeguards (Arts. 16 et seq. FADP)

Our core persistent storage is exclusively in Switzerland. External AI inference takes place in the European Union as described in Section 7.

For every cross-border transfer, we apply the safeguards required by Arts. 16 and 19 para. 4 FADP and Chapter V GDPR:

Adequacy decisions (Art. 16 para. 1 FADP / Annex 1 ODP; Art. 45 GDPR): Transfers to Member States of the European Union and the European Economic Area rely on recognized adequacy determinations.
Standard contractual clauses with swiss adaptations (Art. 16 para. 2 lit. d FADP; Art. 46 GDPR): Our reliance on these safeguards is limited to transfers of personal data necessary **for credit card payment processing and administrative email delivery**. Where transfers to providers of these services are not covered by an applicable adequacy decision, we use the EU Standard Contractual Clauses (Commission Implementing Decision (EU) 2021/914), as recognised by the Swiss Federal Data Protection and Information Commissioner (FDPIC), with the adaptations required under Swiss data protection law and any supplementary safeguards necessary for the relevant transfer.
Swiss–U.S. and EU-U.S. data privacy framework: Our reliance on the applicable Data Privacy Framework is limited to transfers to **U.S.- based credit card payment and administrative email providers** (such as Stripe, Inc.) holding an active certification under the relevant framework that covers the data transferred. We do not treat the United States as generally adequate for all recipients or transfers.

11. Retention and deletion

11.1. Saved tasks, customer content and transformation caches (Processor)

Saved tasks, uploaded files, histories and outputs remain available on our Swiss servers for the agreed service period unless the Customer deletes them earlier or selects a shorter retention setting:

Manual deletion ("Delete" / "Delete all") and "Delete after 24 hours" setting: The Customer can use the "Delete" and "Delete all" controls at any time or enable the optional **"Delete after 24 hours"** setting for individual tasks. The "Delete after 24 hours" setting automatically deletes the selected task and all associated content **24 hours after the task is created**. Manual deletion removes the selected task content from our active Swiss server systems immediately and from the backup servers within 24 hours.

11.2. Account and business records (Controller)

We retain account data for the duration of the contractual relationship and thereafter only as needed to close the account, fulfill statutory obligations or address specific legal claims. Accounting records, invoices, and supporting

documents subject to Art. 958f of the Swiss Code of Obligations are retained for **10 years** from the end of the relevant financial year.

12. Data security, encryption and human access

12.1. Technical and organizational measures (Art. 8 FADP)

We implement technical and organizational measures appropriate to the risks and sensitivity of legal data, including:

Encryption at rest: All Customer Content stored on our physical servers in Switzerland (including files, conversations, databases, retrieval indexes, embeddings, and caches) is symmetrically encrypted at rest using standard encryption methods.
Encryption in transit: All transmissions between client systems, our Swiss gateway, Legislator servers and EU model providers are protected via standard TLS encryption.
Application and boundary controls: Path-traversal prevention middleware, strict input validation without logging request bodies, bot filtering, append-only access audit logs and a strict server-side tool execution sandbox (preventing unauthorized tool calls, shell access, subprocess execution, or arbitrary network access).
Fail-Closed safeguards and incident response: If an approved provider exemption is revoked or an unapproved route is detected, the affected feature immediately fails closed without rerouting Customer Content elsewhere. In the event of a suspected security or confidentiality incident affecting Customer Content, our Head of IT immediately isolates affected services and **notifies the affected law firm's designated contact within 24 hours of awareness** (providing timestamps, affected task IDs, containment steps, and confirmation of encryption integrity), followed by a root-cause and corrective-action report **within 72 hours**.

12.2. Strict restrictions on human access

Customer Content is processed automatically in the ordinary course. Human access to diagnostic or support information occurs only where strictly necessary to resolve a Customer-initiated support request, remediate a technical fault or security incident within the Customer's instructions (via the anonymized procedure in Section 5). All authorized personnel are bound by written, continuing confidentiality duties and professional secrecy as auxiliaries under Art. 321 SCC.

13. Professional secrecy and lawful instructions

The Services are designed to support professional legal work. Where Legislator acts as an auxiliary to a lawyer, notary, or other professional subject to statutory secrecy (Art. 321 SCC, Art. 13 BGFA), our DPA and Professional Secrecy Addendum govern those obligations across the service chain. We remain responsible for our own legal and contractual obligations. The Customer is responsible for ensuring that it is entitled to submit Customer Content, for providing required privacy notices to affected individuals (Art. 19 FADP), and for assessing matter-specific confidentiality restrictions. While engaging a properly bound auxiliary under Art. 321 SCC and Art. 9 FADP does not generally require prior consent from each individual client, the Customer must assess whether specific client Non-Disclosure Agreements (NDAs), contractual cloud/AI prohibitions, or official secrecy duties (Art. 320 SCC) require express client authorization or restrict the use of optional external features (such as external web research under Section 8).

14. Cookies and similar technologies

We use strictly necessary cookies and equivalent browser storage exclusively for authentication, session management, security, and essential platform functionality. We do not use third-party advertising or analytics tracking cookies. If we introduce optional analytics or additional cookies in the future, we will update this Policy and obtain prior consent where required.

15. Automated processing and verification of AI Outputs

The Services use automated processing and Large Language Models (LLMs) to generate legal research results and drafts. Legislator does not use Customer Content to make automated individual decisions producing legal or similarly significant effects on individuals within the meaning of Art. 21 FADP or Art. 22 GDPR.

AI-generated outputs are assistive tools for the Customer's professional assessment and may contain errors, omissions, or fabricated citations ("hallucinations"). In accordance with Art. 12 lit. a BGFA and the SAV/FSA AI Guidelines (Sections 2.2–2.3), **the Customer remains solely responsible for reviewing and verifying all AI-generated outputs against original legal sources before relying on them**.

16. Individual rights and requests

Under the FADP (Arts. 25–29 and 32 FADP) and, where applicable, the GDPR (Arts. 15–22 GDPR), individuals have —subject to statutory conditions and exceptions - the right to request access and information regarding the processing of their personal data, rectification of inaccurate data, erasure ("right to be forgotten") or cessation of unlawful processing, restriction of processing, data delivery or portability (Art. 28 FADP; Art. 20 GDPR), and the right to object to processing based on legitimate interests or direct marketing, as well as to withdraw consent at any time with effect for the future.

For processing for which Legislator AI AG is controller, contact **legal@legislator.ch**. We respond within the statutory time limits: normally within 30 days under Art. 25 para. 7 FADP, or within one month under Art. 12 GDPR (subject to permitted extensions).

For personal data in Customer Content (where we act as processor and auxiliary), requests must be addressed to the responsible Customer (e.g., the law firm). If you contact us directly regarding Customer Content, we will assist or refer the request to the Customer in accordance with the DPA while strictly respecting professional secrecy (Art. 321 SCC, Art. 13 BGFA); we never independently disclose or erase a Customer's legal file in response to an unverified third-party request.

You may raise a concern with the Swiss Federal Data Protection and Information Commissioner (FDPIC) at https://www.edoeb.admin.ch or, where the GDPR applies, lodge a complaint with a competent EU supervisory authority.

17. External websites and independent services

Legal citations, search results, and other links may lead to external websites operated by third parties. If you follow those links or independently use a third-party service, that provider processes data under its own privacy terms.

18. Changes to this policy

We update this Policy when our processing, Services, or applicable legal requirements change. The current version is always available at https://legislator.ch/privacy and replaces all previous versions.

The English version of this Policy is the authoritative version and shall prevail in the event of any discrepancy or inconsistency with a translation into another language.

Annex 1 – Service Providers, Subprocessors, and Processing Locations

This Annex forms an integral part of the Legislator Privacy Policy.

1. Legislator AI AG (Core Platform Operator – Own Infrastructure)

Address: Kirchstrasse 6, 8807 Freienbach SZ, Switzerland (CHE-476.609.141).
Service & Data: Core platform, account database, stored Customer Content, transformation caches, Swiss legal reference database, local search, embedding and reranking models, and daily backups.
Processing Location: Switzerland (proprietary physical servers owned and controlled by Legislator AI AG in Switzerland).
Conditions & Safeguards: Own physical infrastructure in Switzerland; encryption at rest; retention and deletion (immediate or 24-hour auto-erase with 24-hour backup purge) under Section 11.

2. Amazon Web Services EMEA SARL (Subprocessor – Web Reverse-Proxy Gateway)

Address: 38 avenue John F. Kennedy, L-1855 Luxembourg.
Service & Data: Web reverse-proxy ingress gateway routing TLS-encrypted prompts and file attachments directly to Legislator's on-premises servers in Switzerland.
Processing Location: Switzerland (AWS Switzerland Region – Zurich eu-central-2).
Conditions & Safeguards: Data passes exclusively in TLS-encrypted form; 0 days retention (no Customer Content stored). Adequacy for Switzerland/EU (Art. 16 para. 1 FADP) and EU Standard Contractual Clauses with Swiss adaptations (Art. 16 para. 2 lit. d FADP).

3. Google Cloud EMEA Limited (Subprocessor – Approved AI Inference, OCR & Audio Transcription)

Address: Velasco, Clanwilliam Place, Dublin 2, Ireland.
Service & Data: Transient AI inference for approved Gemini and Anthropic Claude models, plus OCR and audio transcription via Gemini API (prompts, relevant document content, generated outputs, and necessary service metadata).
Processing Location: European Union (approved regional EU data centers in EU Member States, including Belgium, the Netherlands, Germany, Finland, and Ireland, with strict regional routing enforcement).
Conditions & Safeguards: Adequacy for EU Member States (Art. 16 para. 1 FADP / Annex 1 ODP). Enterprise Zero Data Retention (ZDR – 0 days content retention) and approved abuse-monitoring prompt-logging exemption (confirmed 1 September 2026 for all Legislator GCP projects); strict no-training prohibition. For any administrative/ support metadata involving non-EEA/Swiss access: EU Standard Contractual Clauses with Swiss Addendum (Art. 16 para. 2 lit. d FADP) and Google LLC's certification under the Swiss–U.S. and EU–U.S. Data Privacy Framework.

4. TensorX Ltd. (Subprocessor – Open-Weights AI Inference)

Address: Unit 25, Classon House, Dundrum Business Park, Dublin 14, Ireland.
Service & Data: Transient AI inference for independently hosted open-weights models, including DeepSeek (prompts, relevant document content, outputs, and service metadata).
Processing Location: European Union (Ireland and Finland) on dedicated sovereign GPU infrastructure.
Conditions & Safeguards: Adequacy for Ireland and Finland (Art. 16 para. 1 FADP / Annex 1 ODP). Zero Data Retention (ZDR – 0 days inference content retention) in volatile memory (RAM) with immediate purge upon request completion; no model training; no data transmission to original DeepSeek developers.

5. Google Cloud EMEA Limited (Subprocessor – Web Research)

Address: Velasco, Clanwilliam Place, Dublin 2, Ireland (with global Google LLC search infrastructure).
Service & Data: Web Grounding for Enterprise. Search queries and contextual information submitted when a user actively selects the optional external web search function (Section 8).
Processing Location: European Union endpoint.
Conditions & Safeguards: Contractual Zero Data Retention (0 days retention of customer data and no model training) under Web Grounding for Enterprise terms; EU Standard Contractual Clauses with Swiss Addendum (Art. 16 para. 2 lit. d FADP) and Google LLC's Swiss–U.S. and EU–U.S. Data Privacy Framework certification for cross-border transfers.

6. Mailgun (Administrative Service Provider – Email Delivery)

Entity: Mailgun Technologies, Inc. / Sinch Email (operating via the EU endpoint api.eu.mailgun.net).
Service & Data: Administrative email delivery containing recipient email address, login code or secure access link, and delivery metadata only; no legal file content, prompts, or documents.
Processing Location: European Union (Germany, Ireland, and Belgium) for message processing; United States for ancillary account/support metadata.
Conditions & Safeguards: Adequacy for EU processing (Art. 16 para. 1 FADP); EU Standard Contractual Clauses with Swiss adaptations (Art. 16 para. 2 lit. d FADP) and Swiss - U.S. Data Privacy - Framework for administrative metadata transfers. Administrative retention under Section 11.

7. Resend – Plus Five Five, Inc. (Administrative Service Provider – Email Delivery)

Address: Plus Five Five, Inc., 2261 Market Street #5039, San Francisco, CA 94114, United States.
Service & Data: Administrative email delivery containing recipient email address, login code or secure access link, and delivery metadata only; no legal file content, prompts, or documents.
Processing Location: United States for email routing and associated administrative metadata storage.
Conditions & Safeguards: Contractual Zero Data Retention (0 days retention for email content). EU Standard Contractual Clauses with Swiss adaptations (Art. 16 para. 2 lit. d FADP). Administrative retention under Section 11 (inference ZDR does not apply to email delivery metadata).

8. Stripe (Administrative Service Provider and Independent Controller – Payments & Subscriptions)

Entities: Stripe Payments Europe, Limited (1 Grand Canal Street Lower, Grand Canal Dock, Dublin, Ireland) and Stripe, Inc. (354 Oyster Point Blvd, South San Francisco, CA 94080, United States).
Service & Data: Billing, payment credentials, transaction references, and fraud-prevention security metadata; no prompts or uploaded legal files.
Processing Location: European Union (Ireland and Germany) and United States.
Conditions & Safeguards: Adequacy for EU processing (Art. 16 para. 1 FADP); Stripe, Inc.'s active certification under the Swiss–U.S. and EU–U.S. Data Privacy Framework, and EU Standard Contractual Clauses with Swiss adaptations (Art. 16 para. 2 lit. d FADP). Statutory 10-year retention for accounting records (Art. 958f CO).

Official Provider Data Protection Links:

Google Cloud: https://cloud.google.com/terms/data-processing-addendum
TensorX: https://tensorx.ai/dpa/ and https://tensorx.ai/privacy-policy/
Amazon Web Services (AWS): https://d1.awsstatic.com/legal/aws-gdpr/AWS_GDPR_DPA.pdf
Mailgun: https://www.mailgun.com/legal/privacy-policy/ and https://www.mailgun.com/legal/dpa/
Resend: https://resend.com/legal/dpa and https://resend.com/security/gdpr
Stripe: https://stripe.com/legal/dpa and https://stripe.com/privacy