Data Processing Agreement

between

[Customer name]
[Address]
CVR no. [ ]
(the "Data Controller")

and

Clerkr ApS
Nørregade 28D, 1.
1165 København K, Denmark
CVR no. 46311809
(the "Data Processor")

(jointly, the "Parties")

Version 2.0, 3 November 2026. This template replaces version 1 of 14 October 2026. The fields for the Data Controller are completed at signature. A Danish version is available at https://clerkr.ai/da/databehandleraftale; in case of discrepancy, the signed version prevails.

1. Background and scope

1.1 The Parties have agreed to this data processing agreement (the "Agreement") to comply with Article 28(3) of Regulation (EU) 2016/679 (the "GDPR") and to ensure the proper protection of the privacy and fundamental rights and freedoms of natural persons.

1.2 The Data Processor develops and provides "Clerkr", an AI work platform for law firms (the "Platform"). The Platform is accessed through a web application and add-ins for Microsoft Word and Outlook, and can be connected to Microsoft Teams and to the Data Controller's document and case management systems. The Platform keeps the material of a matter together, prepares drafts, summaries, tasks and overviews, and answers questions about the matter with references to the source documents, using retrieval-augmented generation: only the most relevant fragments of the Data Controller's own documents are sent, together with the user's request, to a language model for a single inference call. All output is draft material for review by the Data Controller's personnel. The Platform cannot send communications on the Data Controller's behalf.

1.3 In connection with the provision, hosting, support, maintenance and further development of the Platform, the Data Processor processes personal data on behalf of the Data Controller in accordance with this Agreement.

1.4 This Agreement takes precedence over any corresponding provisions in other agreements between the Parties concerning the processing of personal data.

1.5 The following annexes form an integral part of this Agreement: Annex A (Information about the processing), Annex B (Sub-processors), Annex C (Instructions and technical and organisational measures) and Annex D (Other matters agreed between the Parties).

1.6 This Agreement and its annexes shall be kept in writing, including electronically, by both Parties. This Agreement does not release the Data Processor from obligations imposed on it directly by the GDPR or other legislation.

2. The Data Controller's rights and obligations

2.1 The Data Controller is responsible for ensuring that the processing of personal data complies with the GDPR (see Article 24), data protection provisions in other EU law or the national law of the Member States, and this Agreement.

2.2 The Data Controller has the right and obligation to make decisions about the purposes and means of the processing and is responsible for ensuring that there is a legal basis for the processing that the Data Processor is instructed to carry out, including, where relevant, for special categories of personal data, data relating to criminal convictions and offences, and Danish civil registration numbers.

3. The Data Processor acts on instructions

3.1 The Data Processor shall process personal data only on documented instructions from the Data Controller, unless required to do so by EU law or the national law of a Member State to which the Data Processor is subject, in which case the Data Processor shall inform the Data Controller of that legal requirement before processing, unless the law prohibits such information on important grounds of public interest. The instructions are set out in Annexes A and C. Subsequent instructions may be given by the Data Controller during the processing, but shall always be documented and kept in writing.

3.2 The Data Processor shall immediately inform the Data Controller if an instruction, in the Data Processor's opinion, infringes the GDPR or other data protection law of the European Union or a Member State.

3.3 The Data Processor shall not process personal data for any purpose other than those specified in Annex A. In particular, the Data Processor shall never use the Data Controller's personal data or content to train, fine-tune, evaluate or improve any language model or other AI system, and shall ensure that no sub-processor does so.

4. Confidentiality

4.1 The Data Processor shall ensure that persons authorised to process personal data have committed themselves to confidentiality or are under an appropriate statutory obligation of confidentiality, and that access is limited to persons who need it to perform their duties.

4.2 The Data Processor acknowledges that the Data Controller is a law firm and that the personal data and content processed under this Agreement may be subject to legal professional privilege and statutory duties of confidentiality. The Data Processor's obligation of confidentiality with respect to such data is not limited in time and survives the termination of this Agreement and of any other agreement between the Parties.

5. Security of processing

5.1 Taking into account the state of the art, the costs of implementation and the nature, scope, context and purposes of processing, as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons, the Data Processor shall implement appropriate technical and organisational measures to ensure a level of security appropriate to the risk, including, as appropriate, pseudonymisation and encryption, the ability to ensure ongoing confidentiality, integrity, availability and resilience, the ability to restore availability and access in a timely manner, and a process for regularly testing, assessing and evaluating the effectiveness of the measures.

5.2 The measures implemented by the Data Processor are set out in Annex C and in the Data Processor's Security Overview & Technical and Organisational Measures (version 2.0), which forms part of Annex C. The Data Processor may update the measures from time to time, provided that the overall level of security is not reduced.

5.3 The Data Processor shall in all cases ensure that: (a) all personal data and content of the Data Controller is stored and processed, including all language model inference and embedding, on infrastructure located within the European Union; (b) personal data and content is encrypted in transit and at rest; (c) access is restricted to the Data Controller's organisation and, within it, follows the access rights set on each matter; (d) the Data Processor's own personnel can only access the Data Controller's environment while the Data Controller has disengaged the support access lock described in Annex C; and (e) the language model provider retains no prompts or outputs after an inference call.

5.4 The Data Processor shall assist the Data Controller in ensuring compliance with Articles 32 to 36 of the GDPR.

6. Use of sub-processors

6.1 The Data Controller hereby grants the Data Processor a general authorisation to engage the sub-processors listed in Annex B. The Data Processor shall maintain the list and make it available to the Data Controller.

6.2 Before engaging a sub-processor, the Data Processor shall carry out appropriate due diligence and enter into a written agreement with the sub-processor imposing the same data protection obligations as those set out in this Agreement, in particular providing sufficient guarantees to implement appropriate technical and organisational measures in such a manner that the processing meets the requirements of the GDPR. The Data Processor shall provide a copy of the relevant sub-processor agreement to the Data Controller on request; commercial terms may be redacted.

6.3 Where a sub-processor fails to fulfil its data protection obligations, the Data Processor remains fully liable to the Data Controller for the performance of the sub-processor's obligations, in accordance with Article 28(4) of the GDPR.

6.4 The Data Processor shall notify the Data Controller in writing of any intended addition or replacement of a sub-processor at least 30 days before the sub-processor processes any personal data of the Data Controller. The notification shall identify the sub-processor, the purpose, the location of processing and, where relevant, the transfer mechanism.

6.5 The Data Controller may object to the intended change on reasonable grounds relating to data protection by written notice within 14 days of receiving the notification. If the Parties cannot resolve the objection within a reasonable period, the Data Processor may elect not to engage the sub-processor or, where the sub-processor is necessary for the provision of the services, either Party may terminate the affected services by written notice without liability for such termination, and the Data Processor shall refund any prepaid fees for the period after termination.

7. Transfers to third countries or international organisations

7.1 The Data Processor shall not transfer personal data to a third country or an international organisation, and shall not permit any sub-processor to do so, except as set out in Annex B or with the prior written authorisation of the Data Controller. The Data Controller hereby authorises the transfers listed in Annex B, section 5, on the conditions stated there.

7.2 Any authorised transfer shall be based on an adequacy decision of the European Commission or other appropriate safeguards under Chapter V of the GDPR (Articles 44 to 50), including the Standard Contractual Clauses adopted by the Commission. The Data Processor shall assess, and document on request, whether supplementary measures are required.

7.3 Content of the Data Controller (documents, matter data, search index, embeddings, chat history and language model inference) shall under no circumstances be transferred to or processed in a third country.

7.4 If the Data Processor or a sub-processor receives a legally binding request from a public authority, including an authority of a third country, for disclosure of personal data processed under this Agreement, the Data Processor shall: (a) notify the Data Controller without undue delay and no later than 24 hours after becoming aware of the request, unless prohibited by law, and in that case use its best efforts to obtain a waiver of the prohibition; (b) challenge the request where there are reasonable grounds to do so and disclose only the minimum amount of data required; (c) immediately suspend and sever the connection to any sub-processor that has been ordered to disclose the Data Controller's content; and (d) require the sub-processor to delete any content of the Data Controller it holds. The Data Processor confirms that, as at the date of this Agreement, it has never received such a request.

8. Assistance to the Data Controller

8.1 Taking into account the nature of the processing, the Data Processor shall assist the Data Controller by appropriate technical and organisational measures, insofar as this is possible, in fulfilling the Data Controller's obligation to respond to requests for exercising the data subject's rights under Chapter III of the GDPR, and in ensuring compliance with the obligations under Articles 32 to 36 of the GDPR, including notification of personal data breaches, data protection impact assessments and prior consultation with the supervisory authority.

8.2 The Data Processor shall provide the Data Controller with the information necessary to respond to a data subject request within 10 working days of the Data Controller's request. The Data Processor shall not respond directly to data subjects without the Data Controller's instruction.

8.3 The Data Processor may charge the Data Controller for assistance on a time-spent basis at its then-current rates, except where the assistance is required as a result of the Data Processor's breach of this Agreement or the GDPR, or concerns a personal data breach, in which case the assistance is free of charge.

9. Notification of personal data breach

9.1 The Data Processor shall notify the Data Controller without undue delay and in any case no later than 24 hours after becoming aware of a personal data breach affecting personal data processed under this Agreement.

9.2 The notification shall, to the extent known, describe the nature of the breach including the categories and approximate number of data subjects and records concerned, the likely consequences of the breach, the measures taken or proposed to address the breach and mitigate its possible adverse effects, and the contact point for further information. Information that is not available at the time of the first notification shall be provided in phases without undue delay.

9.3 The Data Processor shall cooperate fully with the Data Controller and any supervisory authority in the investigation, mitigation and remediation of the breach, and shall provide the information the Data Controller needs to notify the supervisory authority within 72 hours and, where required, the data subjects.

10. Deletion and return of personal data

10.1 The Data Controller may at any time export its documents in DOCX or PDF format through the Platform and may delete documents and chat history through the Platform. Deletion of a document deletes its derived text fragments and embeddings.

10.2 Upon termination of the services, the Data Processor shall, at the choice of the Data Controller, delete or return all personal data and delete existing copies within 30 days, unless EU or Member State law requires storage of the personal data. On request, the Data Processor shall provide a SQL dump of the Data Controller's organisational data before deletion.

10.3 Backups containing the Data Controller's personal data are retained for a maximum of 30 days and expire automatically. All personal data of the Data Controller is therefore removed from the Data Processor's systems, including backups, no later than 60 days after termination or a deletion request.

10.4 The Data Processor shall confirm deletion in writing on request.

11. Audit, including inspection

11.1 The Data Processor shall make available to the Data Controller all information necessary to demonstrate compliance with Article 28 of the GDPR and this Agreement, and shall allow for and contribute to audits, including inspections, conducted by the Data Controller or another auditor mandated by the Data Controller.

11.2 The Data Processor shall provide the Data Controller with its security documentation, the certifications and audit reports of its sub-processors that it is permitted to share, and a summary of any penetration test on request.

11.3 Audits and inspections shall be conducted upon reasonable prior notice of at least 14 days, no more than once per year unless a breach or a supervisory authority gives cause, during normal business hours, and in a manner that minimises disruption to the Data Processor's operations. The Data Controller shall bear its own costs and the Data Processor's reasonable documented costs of supporting the audit, unless the audit reveals a material breach of this Agreement by the Data Processor.

12. Other matters, entry into force and termination

12.1 The Parties may agree in writing on additional matters not covered by this Agreement; any such agreement shall be set out in Annex D.

12.2 This Agreement enters into force on the date it is signed by both Parties and continues for as long as the Data Processor processes personal data on behalf of the Data Controller. Clauses 4.2 and 10 survive termination.

12.3 This Agreement is governed by Danish law. Disputes shall be settled by the Danish courts.

13. Contact persons

Data Controller	Data Processor
Name: [ ] Email: [ ]	Name: Niels Høgh Høiberg Email: privacy@clerkr.ai Clerkr ApS, Nørregade 28D, 1., 1165 København K

14. Signatures

This Agreement is signed digitally.

For the Data Controller	For the Data Processor, Clerkr ApS
Name: Title: Date:	Name: Title: Date:

Annex A: Information about the processing

Item	Details
Purpose of processing	Provision, hosting, operation, support and maintenance of the Clerkr Platform for the Data Controller, including: storage of documents and matter material; text extraction and chunking; vector embedding for retrieval; retrieval of relevant fragments in response to user requests; generation of drafts, summaries, tasks and edits; connection to the Data Controller's email, Teams and document or case management systems; versioning and audit logging; export; and customer support.
Nature of processing	Automated processing, including collection from connected systems, storage, organisation, retrieval, analysis, generation of draft text, and deletion. Language model inference is stateless: fragments and the request are sent to the model for a single call and are not retained by the model provider.
Types of personal data	All personal data contained in the Data Controller's matter material, including names, contact details, identification numbers (including Danish civil registration numbers), financial information, employment information, correspondence and any other information in documents, emails and messages; and user account data of the Data Controller's personnel (name, work email, role).
Special categories and other sensitive data	Depending on the Data Controller's practice areas, matter material may contain special categories of personal data under Article 9 of the GDPR (for example health data in employment or personal injury matters), personal data relating to criminal convictions and offences under Article 10 of the GDPR and section 8 of the Danish Data Protection Act, and civil registration numbers under section 11 of the Danish Data Protection Act. The Data Processor processes such data solely as part of the matter material, on the Data Controller's instructions and with the measures in Annex C.
Categories of data subjects	The Data Controller's clients and their contact persons, counterparties and their representatives, witnesses, employees of the Data Controller (users) and other natural persons mentioned in matter material.
Duration of processing	For the duration of the Data Controller's subscription, followed by deletion within 30 days and expiry of backups within a further 30 days, in accordance with clause 10.
Location of processing	European Union: Railway (EU region, Netherlands) and Amazon Web Services eu-central-1 (Frankfurt) with EU-only inference profiles. See Annex C, section C.3, and Annex B.

Annex B: Sub-processors

The Data Controller has approved the use of the following sub-processors. The entity names, data categories, transfer mechanisms and links to security documentation are set out in Clerkr's List of Sub-processors, version 2.0, 3 November 2026, which is incorporated in this Annex.

Sub-processor	Purpose	Location	Data categories
Amazon Web Services EMEA SARL	S3 object storage; Bedrock language model inference (Anthropic Claude models)	EU, eu-central-1 (Frankfurt), EU-only inference profiles	Document files; request text and retrieved fragments at inference time (not retained)
Railway Corp.	Application hosting, PostgreSQL database, text extraction, vector embeddings (Clerkr's own model), PDF rendering, LLM observability metadata	EU region (europe-west4, Netherlands)	All structured customer data: users, documents (metadata and parsed text), fragments, embeddings, chat history, tasks, audit logs
PostHog Inc. (EU Cloud)	Product analytics and error tracking	EU (Frankfurt)	Event metadata only; no document content
Resend, Inc. (EU region)	Transactional system email (verification, password reset, notifications)	EU (Ireland region)	Recipient email address; subject and body of system emails; never document content
Microsoft Corporation (only if the Data Controller connects Microsoft 365)	Sign-in, Outlook, Teams and document system connection via OAuth	The Data Controller's own Microsoft tenant	OAuth tokens (encrypted in Clerkr); emails and documents fetched from the Data Controller's tenant
Tavily, Inc. (only if the Data Controller's administrator enables web search; off by default)	Web search to public legal sources	EU (Tavily's EU endpoint)	Search queries typed by the user; never document content

Model provider. Anthropic, PBC is the vendor of the Claude models. The models run inside AWS Bedrock in Frankfurt, operated by AWS. Anthropic has no access to the Data Controller's data and is not a sub-processor. The Data Processor does not use Google Cloud, OpenAI or the Anthropic API directly.

Section 5. Third-country transfers and safeguards. No transfer of the Data Controller's personal data to a third country takes place: all sub-processors process the data in EU regions. Railway Corp., Resend, Inc. and Tavily, Inc. are headquartered in the United States. Their data processing agreements with the Data Processor include the EU Standard Contractual Clauses as a safeguard for any remote support access from outside the EU, and the Data Controller authorises such access on that basis. No other transfer is authorised.

Annex C: Instructions and technical and organisational measures

C.1 Instructions for processing

The Data Processor shall process personal data only in accordance with the instructions in this Annex and any subsequent written instructions from the Data Controller. Specifically, the Data Processor shall: (a) store, host and manage the Data Controller's documents and matter material; (b) collect email, messages and documents from the systems the Data Controller connects, and for email store on the matter only messages that match a matter on an unambiguous signal, keeping for unmatched messages only subject, sender, date and an attachment flag in the user's triage queue; (c) parse and chunk documents into searchable text fragments and create vector embeddings for retrieval; (d) retrieve relevant fragments in response to user requests; (e) generate drafts, summaries, tasks and edits as instructed by users or started by the Platform for connected matters, always as draft material for human approval; (f) keep versions and an audit log; (g) provide customer support; (h) NOT send any communication on the Data Controller's behalf; (i) NOT use the Data Controller's data for training, fine-tuning or evaluation of any model; (j) NOT use the Data Controller's data for cross-customer analytics or search; and (k) NOT disclose the Data Controller's data to third parties except as authorised in Annex B.

C.2 Data Controller access and control

The Data Controller maintains full access to all documents and data within its organisational scope, may export documents in DOCX and PDF at any time, may delete documents and chat history at any time, and may request deletion of any or all data, which the Data Processor shall execute within 30 days. All database queries are scoped by organisation at the SQL layer.

C.3 Data residency and the AWS Bedrock setup

All systems processing or storing the Data Controller's content are hosted within the European Union: web application, workers, database, text extraction and PDF rendering at Railway (EU region); document files at AWS S3 (eu-central-1, Frankfurt); language model inference at AWS Bedrock (eu-central-1 endpoint, EU-only inference profiles); vector embeddings with the Data Processor's own model inside its Railway project (EU); LLM observability metadata on a self-hosted Langfuse instance in the same project; product analytics at PostHog EU Cloud (Frankfurt). Inference never leaves AWS' EU regions. Under the AWS service terms for Amazon Bedrock, inputs and outputs are not shared with model providers and are not used to train models. Amazon Bedrock does not store prompts or outputs, and the Data Processor has not enabled model invocation logging. No content of the Data Controller is transferred to the United States or any other third country.

C.4 Authentication and access control

Email and password sign-in with bcrypt-hashed passwords (cost factor 10); plaintext passwords never logged or persisted; email verification before activation; sessions managed by Better-Auth with HttpOnly and Secure cookies, 30-day validity and 24-hour refresh on activity; single-use, time-limited links (24 hours) for password resets and verification; one organisation and one role per user; access within the organisation follows the rights on each matter; cross-organisation access structurally impossible. Multi-factor authentication (a one-time code sent by email at sign-in, which the Data Controller's administrator can require for all users) is released in November 2026; single sign-on via Microsoft Entra ID is on the 2027 roadmap.

C.5 Support access lock

The Data Controller's organisation has a support access lock in its settings, engaged by default. While engaged, no personnel of the Data Processor can access the Data Controller's content (documents, chat, email and matter material), and support sessions cannot be started. The Data

Controller's administrator may disengage the lock to allow support and engage it again afterwards. Unlocking, locking and every support session are recorded in the Data Controller's activity log, the lock state is visible to its administrators, and production database access by the Data Processor's personnel is logged.

C.6 Boundaries of the AI agent

The Platform has no capability to send email or messages; all outgoing communication it prepares is placed as a draft. Actions with effect outside the Platform (creating a client, moving a file, writing back to the Data Controller's document system) require explicit approval by a user. Every change the Platform makes to a document creates a new version, and the previous version is kept. Sharing a document or sending it for signature is always an action a user takes. Every agent run is logged with matter, instruction, data read and output produced. The agent uses a closed tool registry with no shell, code execution, email or messaging tool. Content retrieved from documents and email is treated as data, not as instructions. The Platform makes no automated decisions with legal effect on natural persons.

C.7 Data isolation and confidentiality

Every database table with customer data carries an organisation identifier injected into every query; the retrieval function includes a non-optional organisation predicate, and all other queries are scoped to the organisation in the application; S3 keys are prefixed by organisation and matter. Product analytics record event metadata only, with session recording and autocapture disabled. Documents are not crawled, indexed externally or used in marketing analytics.

C.8 Encryption

TLS 1.2 or higher in transit on all customer-facing and provider-facing connections, HSTS enforced; AES-256 server-side encryption of document files in S3, with customer-managed keys available on enterprise contracts; disk-level encryption of the database; AES-256-GCM application-side encryption of integration tokens and secrets with the key held only as a production environment variable; backups inherit the encryption of the respective platform.

C.9 Logging and audit

Pipeline logs per processing run; chat history, tool calls, model and prompt version per organisation; retrieval logs with user, organisation and result counts (not content); agent run logs; document version history with author and time; audit log of production database access by the Data Processor's personnel. LLM observability on a self-hosted Langfuse instance in the EU records timing, model and token counts only, never prompts, outputs or content. No third-party error tracking or external LLM observability service is used.

C.10 Access by the Data Processor's personnel

Production secrets held only in the production environment; a short list of named engineers with production access, revoked on departure; MFA on all cloud accounts and SSO on Railway; a dedicated IAM user scoped to the Clerkr S3 buckets and Bedrock invocation; no database access from outside the private network; developer access to production data only through an explicit, time-limited session, subject to the support access lock, and logged; no administrative interface exposing raw personal data.

C.11 Supply chain, change management, vulnerability management

Dependencies pinned via lockfile; builds from the lockfile only; pull request and CI (lint, type check, schema validation) before production; separate development, staging and production environments; no customer production data in development or test; security advisories reviewed and patched in the regular release cycle, critical vulnerabilities out of cycle; external penetration test scheduled before

the end of Q1 2027 and at least annually thereafter, with a summary available to the Data Controller on request.

C.12 Backup and disaster recovery

Daily database backups and S3 object versioning; backup retention 30 days; backup systems with the same security measures as production; Recovery Point Objective 24 hours and Recovery Time Objective 8 business hours.

C.13 Breach notification and cooperation

Platform and pipeline failures logged and reviewed by the Data Processor's engineers, with automated alerting being added; notification of the Data Controller without undue delay and no later than 24 hours after becoming aware of a personal data breach, with the information in clause 9; full cooperation in investigation and remediation; assistance with the Data Controller's notification to the supervisory authority and data subjects.

C.14 Data subject rights

Assistance with access, rectification, erasure, restriction and portability; information provided within 10 working days of the Data Controller's request; no direct response to data subjects without instruction. Rectification is made in the source document, after which fragments and embeddings are regenerated. Deletion of a document deletes its fragments and embeddings.

C.15 Retention

Customer content: duration of the subscription, deletion within 30 days of termination or request. Backups: 30 days. User account data: deleted within 90 days of termination. Security and operational logs and analytics events: 12 months. LLM observability metadata: deleted with the organisation. No retention at AWS Bedrock.

C.16 Third-country transfers and government requests

No personal data of the Data Controller is transferred to a third country; all sub-processors process in EU regions, including system email (Resend, EU region) and, if enabled, web search (Tavily, EU hosting). Providers headquartered outside the EU have the EU Standard Contractual Clauses in their data processing agreements as a safeguard for remote support access, as set out in Annex B, section 5. AWS has committed to notify customers of government requests for customer data where legally permitted and to challenge inappropriate requests. The Data Processor applies the procedure in clause 7.4, including immediate severing of the connection to any sub-processor ordered to disclose the Data Controller's content.

C.17 Organisation and personnel

Written information security and data protection policies with defined roles; confidentiality obligations and annual training for all personnel with access to customer data; assessment of sub-processors before engagement and flow-down of equivalent data protection obligations; named security and privacy contact (clause 13).

Annex D: Other matters agreed between the Parties

None.