CLEVER
CONTRACT INTELLIGENCE

17 October 2026 · fifth edition

Data security and compliance

This document states where contract data is held, who sees it, where it is sent for analysis, and what CLEVER explicitly does not provide. Each statement names the standard, that is the shared platform with separated tenants, and the option with which a customer goes beyond it on request. Options apply by agreement and are not presented as available where they are not.

Item	Standard (shared platform)	On request (by agreement)
Application	delivered via Netlify, holds no contract data	unchanged
Database	one Supabase project, region Frankfurt (AWS eu-central-1), since 19 September 2026; previously Ireland	own project per customer, EU region of the customer's choice
Separation between customers	per tenant by a database rule on 39 tables (decision of 21 September 2026); test on a copy of the schema: no foreign row visible, in five roles and two tenants	own database: no query path between customers
Access within the tenant	governed by role (database rules per role); a deal manager sees only the contracts assigned to them	unchanged
Analysis by language model	Anthropic (contracting party: Anthropic Ireland, Limited), processing outside the EEA; inputs and outputs deleted within 30 days	processing in the EU (Amazon Bedrock, Frankfurt) as a bookable option; switched on for the customer who books it
Model training	excluded	excluded
Operator access	technically possible via the database, occasion-based; the members of a tenant are set by its administrator, not by the operator	only once enabled by the customer
Assessment logic	part of the delivered application, no customer data	server-side with an own instance
Source code	private repository	unchanged; with an own instance the customer receives a versioned release
Encryption	in transit and at rest	unchanged

Certification	none	none
Data processing agreement	on request	part of the contract

— Where the platform runs

The application is delivered via Netlify and holds no contract data. It loads that data from the database only after sign-in. Database and sign-in are hosted by Supabase in the Frankfurt region (AWS eu-central-1), within the European Union.

Until 19 September 2026 the database was held in the Ireland region. The holding has moved to Frankfurt in full. The former project in Ireland is paused, no longer in use, and keeps the state of the day of the move as a fallback.

Transmission to the language model passes through a server function at Netlify. It forwards the text and logs token counts and durations, not contract content.

Which files the website serves is governed by an allow list. Every folder is either blocked or explicitly released; the check runs automatically on every change. On 17 September 2026 it found four internal folders without customer data that were retrievable. They are blocked.

— Tenants and own instance

Customers work as separate tenants in a shared database. Each of the 39 tables holding customer data carries a tenant key, and an additional rule lets an account see only the rows of its active tenant. The existing role logic applies on top, unchanged. The database sets the tenant key itself and rejects a foreign value sent by the browser. An account can be a member of several tenants; which members a tenant has is decided by its administrator. The operator creates tenants (with a first administrator) and can shut them down; it adds no members and assigns no roles. In the application it sees no tenant content; technical access through the database is separate and described in the section on operator access.

Test — on 22 September 2026 the separation was tested on a copy of the database schema with all 103 rules: across 39 tables, five roles and two tenants no foreign row was visible (25 of 25 checks), and the existing role logic remained unchanged (195 of 195 pairs of role and table). This is a test on a copy, not an audit by a third party.

Residual risk — one rule separates the customers; an error in it would affect all tenants of the same database. Against it stand the additional rule per table, the setting of the key by the database and the repeatable test. Whoever needs hard separation receives an own instance on request: own database, own credentials, optionally in the customer’s own cloud account, by agreement. Today it exists as an offer, not as a prepared environment.

— Who sees uploaded contracts

Within a customer: administrator, risk manager and approver see all contracts. A deal manager sees those assigned to them, the buyer role the contracts of the buying side. The database has enforced this since 5 August 2026, including on direct access to the interface.

Outside the customer: Anthropic for analysis, see below. The operator currently has technical access to the database for operation and troubleshooting.

Stored are the contract text read from the document, the file name and the analysis. The original file is not stored.

Questions to the platform assistant are stored and sent to the operator as a list by e-mail once a day. The sending service is Resend, based in the United States.

— Processing by the language model

Analysis currently runs through the Anthropic interface (api.anthropic.com). The contract text is transmitted for that purpose and leaves the European Economic Area; on this route the region of processing is not fixed to the EU. The basis is Anthropic's commercial terms with the data processing addendum and EU standard contractual clauses. For customers in the EEA the contracting party is Anthropic Ireland, Limited (Commercial Terms of Service, definition of "Anthropic"). Anthropic publishes its sub-processor list at anthropic.com/subprocessors.

No training — Anthropic does not train on this content; its commercial terms exclude it (Commercial Terms of Service, section B). CLEVER does not train any model itself. The interface key is held on the server and never reaches the browser.

Retention — Anthropic deletes API inputs and outputs within 30 days, except where required by law or for usage policy enforcement (Anthropic Privacy Center, "How long do you store my organization's data?"). The provider's standard retention applies; no deviating agreement exists.

Processing in the EU — CLEVER offers analysis through Amazon Bedrock in the Frankfurt region as a bookable option (module "Processing in the EU", M28). It is switched on for the customer who books it. CLEVER does not use this route itself, and no analysis runs through it today. According to AWS documentation, Bedrock does not store inputs and outputs by default, and model providers have no access to them. On this route AWS acts as processor.

— Assessment logic and source code

The assessment logic comprises the instructions to the language model and the risk registers. Today it is part of the delivered application and contains neither customer data nor access keys. With an own instance it is held server-side.

The source code is held in a private repository accessible only to the operator. Only the built application is served. With an own instance the customer receives a versioned release.

— Further services

Resend sends the platform's notifications, sender noreply@clever-ai.org. Sign-in is by sign-in link or with a Google account; sign-in with a Microsoft account is built in, its set-up for the database in Frankfurt is still pending. E-mail to people at clever-ai.org runs through Microsoft 365. A payment provider is prepared but not in operation; it will process payment data, not contract data. The interface loads typefaces from Google Fonts, which transmits the visitor's IP address to Google. Google Drive is contacted only when a user inserts a link to a document stored there.

The interest form on the start page accepts submissions only within limits: a hidden decoy field, a minimum of 2.5 seconds on the form, at most two submissions per address and hour, five per connection and hour and twenty per day, and a limit on links in the text. Counting runs in the memory of the server function and is therefore not complete protection.

— Log and deletion

Deletions, restorations, permanent deletions and compliance holds are recorded in a log that accepts entries only. A deleted contract remains restorable at first; only permanent deletion removes it together with all risk positions. On request we remove a holding in full.

— What we do not claim

CLEVER is not certified. There is no ISO 27001 certification and no off-the-shelf data processing agreement for customers. Separation between customers rests on a database rule in a shared database; it is tested on a copy, not by a third party. Separate databases exist only as an option by agreement, not as the standard.

CLEVER does not claim that analysis runs in the EU today. Storage is in Frankfurt; analysis runs through Anthropic outside the EEA. Processing in the EU is a bookable option and is switched on only for the customer who books it.

Corrections — the second edition of this document named Frankfurt as the storage location while the database was still in Ireland, and stated that contract texts do not leave the EEA and are not stored by the model provider. The third edition corrected all three statements. Since 19 September 2026 the Frankfurt location is accurate; the other two statements remain inaccurate for today’s analysis route. The third edition stated retention of up to 30 days at Anthropic as fact without a source; the fourth edition marked it as an assumption. Since 22 September 2026 it is documented with a source. The fourth edition named separate databases per customer as the target architecture and stated that separation by database rule does not exist. With tenant separation in operation it is the other way round: the database rule is the standard, an own database is an option.

Questions: Dr. Jürgen Tiedge, Managing Director of Kahlstorf GmbH (CLEVER) · clever-ai.org