Eigenwelt Plus · Eigenwelt Labs

Terms & Conditions

Last updated 17 November 2026

Eigenwelt Plus is the hosted inference and team platform operated by Eigenwelt Labs UG (haftungsbeschränkt), Berlin, Germany ("Eigenwelt", "we", "us"). These Terms govern the platform at platform.eigenweltlabs.com, paid Eigenwelt model inference, and connected team, Knowledge Hub, intake, and task features, and managed connector sign-in (the "Service").

The Service is for businesses and professionals. If you accept these Terms for a law firm or other organisation ("Customer"), you confirm that you may bind it. The free LegalWork desktop App has separate terms.

The service

Eigenwelt Plus provides organisation accounts, paid model inference, premium features, member and settings management, usage reporting, a Knowledge Hub for supported workflows, integrations, plugins, MCP servers, and presets, and intake, which turns incoming emails and API submissions into assigned tasks. Custom deployments, customer-specific models, and continual-learning work may use a separate order form.

There are three plans, Plus, Pro and Max. All include the Eigenwelt models with the included usage described below, plus the Knowledge Hub and intake; they differ in price and in the amount of included AI usage. The plan, price, and scope are those shown in the order and on the billing page.

Platform hosting

The platform hosting region is selected when the firm is set up or agreed with Customer: EU on Scaleway in Paris (fr-par), or US on Google Cloud in Iowa (us-central1). It determines the location of the platform, gateways, databases, and stored files and applies to all members. Selecting the US platform is an instruction to process platform data there; the DPA requirements for international transfers continue to apply. A later change of platform hosting region requires a separate agreement. Model regions are configurable independently. Shared services and their processing locations are described in the Privacy Notice and DPA.

Managed connector sign-in

As an optional feature, the platform brokers sign-in to supported Google and Microsoft accounts through Nango Inc., United States. Nango stores and refreshes OAuth credentials; Eigenwelt verifies that the connection belongs to the user and firm and provides access tokens to the LegalWork server. Email, calendar, file, and attachment content itself travels directly between LegalWork and the connected provider. US processing of authentication and connection data is independent of model-region settings and is not covered by ZDR for inference content.

Customer ensures that users may connect the accounts and grant the required permissions, and that required notices and approvals are in place. The connected provider's terms govern its service. Disabling write actions in LegalWork does not reduce OAuth permissions already granted. Users can disconnect and revoke permissions at the provider. Provider or Nango restrictions may interrupt a connection or

require reconnecting. An own-app OAuth connection authorizes directly with the provider and does not need this authentication service.

Data processing and international transfers are described in the Privacy Notice and the data processing agreement. Changes to subprocessors follow the notice and objection procedure set out there.

Model inference, regions, and ZDR

Eigenwelt's gateway does not store prompt or output content and retains only content-free metering data such as model, token counts, cost, status, and timestamps. We do not use paid inference content to train shared or public models.

Every model carries a processing region, shown in the platform and in LegalWork. The initial model region matches the firm’s platform hosting region; model access also depends on the plan. Mistral’s EU regional endpoint can also process requests in EFTA countries. Where “All models” is enabled for an allowed region, this includes models added there later. With individual model selection, new models stay off until an administrator enables them. Subprocessor changes remain subject to the DPA notice and objection procedure.

We obtain Mistral Large 4 (Eigenwelt Europe Mistral) directly from Mistral AI through its EU regional endpoint. ZDR is enabled for Eigenwelt; use of API data for training is disabled. Details are in the Privacy Notice and data processing agreement.

An administrator can enable further regions, currently the United States. That selection applies to all of Customer's users and is a documented instruction to send inference content to model infrastructure in the selected region. Such transfers are covered by an adequacy decision or appropriate safeguards under Chapter V GDPR. Customer is responsible for assessing whether a region is consistent with its confidentiality, professional, and data-protection duties.

Zero data retention (“ZDR”) is enabled for Eigenwelt’s accounts with our inference subprocessors; prompt or output content is not used to train models. The provider-specific retention exceptions described in the Privacy Notice and data processing agreement apply, for example where required by law or to address misuse. ZDR does not apply to account, billing, security, and usage records or to Knowledge Hub items Customer intentionally stores. Per-provider details and the list of providers are in the Eigenwelt Plus Privacy Notice.

ZDR also applies to the Gemini models served through Google Cloud (Vertex AI): Google has switched off abuse-monitoring logging for Eigenwelt, so prompts sent to these models are not logged or retained for abuse monitoring either.

ZDR also applies to the Claude and GPT models. We obtain them through OpenRouter's regional hosts only from approved zero-data-retention cloud endpoints at Microsoft Azure, Amazon Web Services, and Google Cloud. For the Claude models, we may also use Amazon Bedrock directly. EU and US inference use separate regional inference profiles within the same AWS account, according to the enabled hosting region. The providers and processing locations are listed in the data processing agreement. Requests are not sent to Anthropic's or OpenAI's own interfaces.

For Claude and GPT models, a technical failure or capacity shortage may cause a request to be routed to the same model version through another approved route, including OpenRouter as a fallback. The same model permissions, selected hosting region, ZDR subject to the provider-specific exceptions described, and displayed model prices apply. EU models may use only model endpoints within the EU; a general classification as “Europe” is insufficient. If no suitable route is available, the request fails.

Eigenwelt Europe EigenJev is operated by Eigenwelt on Google Cloud in the Netherlands (europe-west4). Eigenwelt US Jev is provided by TypeSafe AI, Inc. in the United States, with ZDR enabled for Eigenwelt's TypeSafe account. An administrator must enable both the US region and the model to use Eigenwelt US Jev.

Accounts and administrators

Each account and seat is for one named user. Users must keep credentials secure and may not share accounts, tokens, or API keys. Customer is responsible for its users and for removing access when no longer authorised.

Administrators may manage users, seats, billing, usage settings, and shared Hub items, and decide which models and hosting regions the firm uses. Those choices apply to all of Customer's users. Customer is responsible for appointing administrators and obtaining required notices and consents. Administrator actions bind Customer.

Plans, seats, and payment

Paid plans are recurring and billed per seat. The plan, price, billing interval, seat count, included usage, and features shown at checkout, in the dashboard, or in an order form are part of this agreement. Prices exclude VAT and other taxes unless stated otherwise; VAT is calculated and shown at checkout.

Each plan is offered on two billing intervals. On yearly billing the term is one year, the fee is due in advance for the year, and the subscription renews for a further year unless cancelled before renewal. On monthly billing the term is one month and renews monthly. The billing interval is chosen at checkout and can later be changed through the payment provider's customer portal where offered; a change takes effect at the next renewal.

Self-service subscriptions renew until an administrator cancels. Cancellation and plan downgrades take effect at the next renewal. Seat additions and upgrades may take effect immediately with charges prorated over the remaining term; on yearly billing an added seat is prorated to the end of the current year. Seat reductions cannot go below current members and pending invitations and do not create a mid-period credit. Paid fees are non-refundable except where required by law or stated at purchase.

Trial

A firm's first eligible subscription may start on Plus, Pro or Max with a seven-day trial for a one-time €1 fee including VAT. The trial fee is charged at checkout, separately from the subscription, and is not credited against later subscription payments. The trial is for one user and includes €30 of AI usage in total, regardless of the selected plan. Inviting additional users requires starting the paid subscription early. Unless cancelled before the trial ends, the selected paid subscription then starts. Cancel the trial in the dashboard under “Billing”. Trials already in progress retain the allowance promised at their checkout.

At most one trial is granted per firm; a prior subscription disqualifies. We notify Customer in text form before the trial ends.

Included AI usage

Each paid seat includes the inference allowance shown for its plan. The allowance is granted monthly, including on yearly billing. The app shows when it next resets. Usage depends on the model and input, cached-input, and output tokens processed. Allowances are assigned per user, do not roll over into the

next usage period, have no cash value, and cannot be transferred or resold. Existing subscriptions retain their current prices, usage limits and access while continuing on their existing offer.

Usage is counted against the allowance, and against any purchased credits, at the model prices in effect when the request is made. Model prices are shown per model for input, cached-input, and output tokens on the platform's Models page. They consist of the price our provider charges for the model plus a platform fee, and follow changes in provider prices, exchange rates, and the platform fee. We may change model prices with immediate effect for future requests; usage already processed is not repriced. Model prices are not part of the plan, and changing them is not a plan change under “Changes, suspension, and termination”.

Requests pause when a user exhausts their included allowance and resume at the displayed reset time. The app shows usage consumed and when the allowance next resets. If optional additional usage is offered and enabled, included usage is consumed first and further usage draws from purchased credits at displayed rates and may be subject to a separate monthly cap. Administrators may be able to set spending caps, depending on their plan. Credits are non-transferable and non-refundable except where law or purchase terms say otherwise. Metering may slightly overshoot a limit because final usage is known only after a request completes.

Knowledge Hub and Customer content

Hub items are shared with authorised members according to roles and permissions. Customer is responsible for what its users publish and must have the right to share it. This includes items the assistant creates or writes into the shared library on a user's instruction, which count as published by that user. Credentials may be included only through an expressly protected sharing feature. Automated secret detection does not replace user review.

Customer retains its rights in prompts, Hub items, configurations, and other content. To the extent permitted by law, Customer owns generated output. Customer grants us and our service providers only the rights needed to provide, secure, and support the Service and comply with law.

Intake and tasks

Intake is off until an administrator turns it on. Customer decides which senders are admitted, who sees tasks, and which model intake uses, and is responsible for informing senders about the processing of their messages where required. Emails sent to intake addresses are received by our provider Brevo; details are in the Eigenwelt Plus Privacy Notice.

AI processing for intake has its own budget for the firm, shown in the intake settings. Once it is reached, new messages still become tasks, but without AI extraction and assignment, until the budget resets. Tasks and assignments created by AI must be reviewed like any other AI output.

Confidentiality

Each party will protect the other's confidential information with reasonable care, use it only for this agreement, and disclose it only to personnel and providers who need it and are bound by confidentiality, or where law requires disclosure. Customer Content is Customer's confidential information.

Acceptable use and professional responsibility

Customer must use the Service lawfully and in line with professional, confidentiality, privilege, supervision, and data-protection duties. Users must not bypass seat, usage, safety, or security controls, interfere with the Service, or violate another person's rights.

Eigenwelt Plus assists with legal work but does not provide legal advice. AI output can be incorrect, incomplete, or outdated. A qualified professional must review and verify output and citations before relying on or sharing them.

Changes, suspension, and termination

We may update models and features or apply reasonable safety, security, rate, and capacity controls. We may suspend access for a security risk, unlawful use, breach, non-payment, or provider outage and will use reasonable efforts to give notice and restore access where possible.

Material adverse plan changes receive at least 30 days' notice and apply no earlier than the next renewal unless required sooner by law or urgent security needs. Compliance with the enabled model regions and ZDR as described above are core paid commitments; we will not change them to Customer's detriment without notice under this paragraph. When a subscription ends, hosted inference and premium features stop. Customer should export Hub data it needs before termination.

Warranty and liability

Except for express commitments above, the Service is provided “as is” and “as available”. We do not warrant uninterrupted service or accurate output. Nothing limits liability for intent, gross negligence, injury to life, body or health, an expressly assumed guarantee, or mandatory liability. For slight negligence, we are liable only for breach of a material contractual obligation and foreseeable damage typical of this agreement.

Changes to terms, law, and contact

We may update these Terms prospectively with reasonable notice of material changes. German law applies, and the courts of Berlin have jurisdiction to the extent permitted by law.

Questions: chris@eigenweltlabs.com or eigenweltlabs.com/contact.
