ZeeDo Privacy Policy

ZeeDo

Version 2 · Last updated 5 November 2026 · Fingerprint 529666f2b89f

Legal texts are published in English; English binds. This version is kept unchanged. Every version, each with its PDF: https://zeedo.cloud/en/legal/versions/privacy

1. Who is responsible

1.1 The controller for the processing described in section 4 is ZeeDo GmbH, c/o Ramzy J A Abushaaban, Elsa-Neumann-Straße 53, 13629 Berlin, Germany, contact@zeedo.cloud (**"ZeeDo", "we", "us").** For questions about data protection, write to privacy@zeedo.cloud.

1.2 We have not appointed a data protection officer. Questions about data protection go to privacy@zeedo.cloud. If we appoint a data protection officer, we will name them here.

1.3 Our competent supervisory authority is the Berliner Beauftragte für Datenschutz und Informationsfreiheit (Berlin Commissioner for Data Protection and Freedom of Information).

1.4 Words our Terms of Service define, such as Organisation, Workspace, Member, Bot, Customer Content, BYOK and ZeeDo's model, have the same meaning here. A "customer" is the business or professional for which an organisation's account is held. "GDPR" means Regulation (EU) 2016/679.

2. Summary

- ZeeDo is a service for businesses and professionals. It is not offered to consumers or to anyone under 18.
- Customer data is stored and processed in the EU (France, Germany and Ireland). Section 8 names the sites and the exceptions.
- Bots use ZeeDo's model unless a customer uses its own key or switches it off. Its requests go only to providers that keep no copy and do not train on them, and for an organisation of our EU region they are processed in the EU, except two checks a model in the United States may make unless an organisation keeps them in the EU: reading the rules people set for a bot, and checking a button a bot is about to press, which sends its labels only (sections 5.2, 7.2 and 8).
- For the data our customers' bots read and create, the customer decides and we act on its instructions (section 3).
- A customer controls how long its bots' memory is kept. Raw memory episodes are kept at most 90 days (section 5.3).
- An erasure request runs one job. Within 72 hours it permanently deletes what the identifiers given find in bot memory, chats, runs, traces and the files people handed bots. Some data it does not search, and some the identifiers cannot find: the customer deletes that data, or it is deleted when its own period ends (section 9.2).
- Audit and consent records are kept for 400 days, also after an erasure request, because they

may be needed to establish, exercise or defend legal claims (section 9.2).

- We do not use Customer Content to train or fine-tune any AI model, and ZeeDo's model uses only providers that do not train on it.

3. Two roles: controller and processor

3.1 **ZeeDo as controller.** We decide how and why we process data about the people who deal with us directly: Members of our customers' organisations, billing contacts, people who write to us or send us a notice or a security report, and subscribers to our status page. Section 4 describes this processing.

3.2 **ZeeDo as processor.** Our customers use bots that read and act on data, including personal data about third parties (for example, emails, tickets, customer records, documents and web pages). For that data, which is part of Customer Content, **the customer is the controller** and we process it on the customer's behalf under our Data Processing Addendum (DPA). The customer is responsible for having a lawful basis for that processing and for informing the people concerned. Section 5 explains what we store in that role, so that people whose data a bot handles can understand it.

3.3 **If a ZeeDo bot handled your data** and you are not a ZeeDo customer, the organisation that runs the bot is responsible for that processing. A ZeeDo bot names that organisation: our systems add the line — `{bot name}, an AI assistant at {org name}` to the first message of each conversation or thread on each channel, and to the signature of every email a bot sends, for example — `Kai, an AI assistant at Trailpost.` Please contact that organisation. If you contact us, we will pass your request to it, as its processor.

4. Data we control

Category	What it includes	Purpose	Legal basis (GDPR)
Account and identity	Name, email address, organisation and role; the method of your last sign-in and of your last confirmation, and when you last used a passkey; your id in your organisation's directory, where it connects one; your notification settings and the browsers you turned push notifications on for. Sign-in, including passkeys, is provided by WorkOS in the United States (8.2). Inside ZeeDo a person is identified by a member id. The WorkOS subject id is stored in one place and is never displayed or exported	Provide the account, sign-in and access control	Art. 6 (1) (b) where you are the customer (a sole professional); otherwise Art. 6 (1) (f): our and the customer's legitimate interest in providing the service the customer ordered to the people it admits

Category	What it includes	Purpose	Legal basis (GDPR)
Signup affirmation and terms acceptance	That you act for business or professional purposes, are 18 or older and, where you accept for an organisation, have authority to bind it; the terms version accepted; the time; the version of the text shown	Prove who may use ZeeDo and which terms apply	Art. 6 (1) (b) and (f): proving eligibility and the terms that apply
Contract confirmations and acceptance records	Who accepted which version of the Terms of Service, the Acceptable Use Policy and the DPA, and when; the confirmations of the contract (4.4) and the addresses they are sent to; objections to changes and acceptances of them	Prove the contract and the terms that apply to it	Art. 6 (1) (c) for the confirmations, which we must keep as business letters (§ 257 HGB, § 147 AO); Art. 6 (1) (f) for the acceptance records: establishing, exercising and defending legal claims
Sessions and security	Session cookie, IP address, user agent, and a hash of your browser's user-agent and language headers, which binds the session to that browser	Keep accounts secure; detect account takeover	Art. 6 (1) (f): keeping accounts and the service secure
Audit and consent records	For each action: a pseudonymous actor id, the action, the resource, hashes of the state before and after, IP address, user agent, trace id, time. For each consent: the id, version and hash of the exact text shown	Prove who authorised what, and when; establish, exercise and defend legal claims	Art. 6 (1) (f): proving who authorised what. They are kept after an erasure request under Art. 17 (3) (e)
Billing	Legal entity name, billing address, VAT ID, billing currency, invoices, subscription items, and the billing emails we send. Stripe collects payment details; we do not store card data	Charge for the service; meet tax and accounting duties	Art. 6 (1) (b) where you are the customer; otherwise Art. 6 (1) (f): dealing with the customer's billing contacts; Art. 6 (1) (c) for invoices and the records tax and commercial law require us to keep. Stripe processes billing data for us as our processor, and processes payments and prevents fraud as an independent controller under its own privacy policy
Usage and fair use	Machine utilisation, working hours, run counts, cost records (model provider, model, token counts, estimated cost), fair-use flags	Show usage to the customer; enforce budgets and fair use; plan capacity	Art. 6 (1) (b) and (f): running the service fairly for every customer

Category	What it includes	Purpose	Legal basis (GDPR)
Abuse and suspension records	The reason for a restriction, the notices sent, appeals and their outcome; notices of illegal content we receive, and the details of the people who send them	Protect the service and third parties; handle notices and appeals	Art. 6 (1) (f): protecting the service, our customers and third parties; Art. 6 (1) (c) where the Digital Services Act requires us to handle a notice or give reasons
One-per-person record of free AI usage	Keyed one-way codes of the sign-in id, the email address and the company domain of the person who created an organisation, and a per-day code of the network it was created from; never the address or the domain itself	Give the free AI usage of the Free plan once per person and per company, and stop it from being claimed again by deleting and re-creating an account	Art. 6 (1) (f): preventing abuse of a free offer
Security disclosure reports	Your report, your contact details, our triage notes	Handle vulnerability reports	Art. 6 (1) (f): securing the service
Support and correspondence	Messages you send to our mailboxes or through our contact form, and our answers	Answer you	Art. 6 (1) (b) where it concerns a contract with you; otherwise Art. 6 (1) (f): answering the people who write to us; Art. 6 (1) (c) for messages we must keep as business letters (§ 257 HGB, § 147 AO)
Website and server logs	IP address, time, requested page, user agent	Operate and secure the website and the service	Art. 6 (1) (f): operating and securing the service
Monitoring and status	Application error reports (the type of error, where in our code it happened and the pattern of the page, and nothing a person or a bot wrote); two page-speed figures with no identifier; uptime checks of our own service; the email address of a status-page subscriber; infrastructure logs, metrics, alerts and traces of our own services' requests, which carry identifiers, times and network addresses (4.2)	Find and fix faults; keep the service available and fast; tell subscribers about incidents	Art. 6 (1) (f): running a reliable and secure service

4.1 **No tracking cookies.** The website and the product set only cookies, and keep only data in your browser's storage, that are strictly necessary to provide the service you ask for (section 25

(2) no. 2 of the Telecommunications Digital Services Data Protection Act, TDDDG): for example the session cookie that keeps you signed in, cookies that keep a form you are filling in or a notice until the next page, the language and appearance you chose, and an id that keeps the live updates of each browser tab apart. We use no advertising or analytics cookies and no tracking, so we do not ask for cookie consent.

4.2 **Monitoring.** We use:

- **Sentry (**Functional Software, Inc.), in its EU region, for application errors. An error report is built from a fixed list of fields: the type of error, where in our code it happened and the pattern of the page, not its address. It holds nothing a person or a bot wrote. Browsers send their reports to our own servers, which pass them on, so Sentry does not receive your network address from them.
- **Our own measurement** of two page-speed figures, Largest Contentful Paint and Interaction to Next Paint. Your browser sends each to our own servers as a name and a number, with no cookie, no identifier and no page address, and we add it to our own performance statistics, which keep no single report.
- **Better Stack,** for uptime checks of our service and for our status page. If you subscribe to status updates, Better Stack holds your email address until you unsubscribe.
- **Grafana Labs (**Grafana Cloud, in an EU region in Germany), for our infrastructure logs, metrics and alerts, and the traces of our own services' requests (7.6).

4.3 **Notifications.** We send alerts and requests for approval to people in the app, by email and by push notification, as each user or workspace chooses, and through no other channel. Email is sent through Amazon SES in Amazon Web Services' Frankfurt region. Push notifications to a browser use the standard Web Push protocol and are sent from our own servers; the message is encrypted end to end and passed on by the push service of your browser's maker, which cannot read it. If we release native mobile apps, their notifications will go through Apple's and Google's push services, and we will add those to the sub-processor list first.

4.4 **Contract confirmations.** When an organisation is created, and again when its first paid plan starts, we will send the person who accepted and the Owners an email confirming the contract. It will name the versions in force of the Terms of Service (with the Acceptable Use Policy), the Data Processing Addendum and this Privacy Policy, give the address at which each version is kept, and attach each as a PDF. We will keep every published version of each of these documents, unchanged, at its own address on zeedo.cloud.

4.5 **What you must give us, and where else data comes from.** To create or join an organisation, you must give us your name and email address through your sign-in, and make the affirmation in the table above. Without them we cannot provide an account. A paid plan also needs billing details. Everything else arises from your use of ZeeDo. Some data we do not collect from you: an Owner or Admin who invites you gives us your email address; your organisation's single sign-on or directory gives us your name and your id in it, through WorkOS; and Stripe tells us about payments.

5. Data we process for our customers (bots, memory and recordings)

This section describes what we store as a processor. The customer controls this data.

5.1 **What a bot may read.** Anything the customer connects it to: mailboxes, documents, tickets, code, customer records, calendars and web pages. In its browser, a bot reads a web page as a structured text outline, or as a picture of the screen in which password and payment fields are masked. The services a customer connects are the customer's own providers, under its own contracts with them, and not our sub-processors.

5.2 **Where it goes.** Customer data is stored and processed in the EU (France, Germany and Ireland): a bot's machine, and the stores that hold Customer Content, among them the database, the workflow history, the object store, the trace store, snapshots and recordings (section 8.1). When a bot calls an AI model:

- with the customer's **own key (BYOK),** the prompt goes to the model provider the customer chose, which acts as the customer's processor under the customer's own contract with it;
- with **ZeeDo's model:** ZeeDo's model is the default for every workspace on every plan, Free included, unless a bot or its workspace uses the customer's own key (BYOK) or an Owner or Admin has switched ZeeDo's model off for the workspace. Its requests go to OpenRouter and to the model provider named for each model on the sub-processor list: each model is served by one named provider. Every request is restricted to providers that keep no copy of prompts or responses (zero data retention) and do not train on them, and we set this restriction on each request ourselves. For an organisation of our EU region, requests go only through OpenRouter's EU endpoint and are processed in the EU; a request that cannot be served there fails and is not sent elsewhere. For an organisation of another region, requests go through the endpoint we name for that region.

Two uses may be served from the United States, by a decision model run by Perplexity and reached through OpenRouter, once ZeeDo uses that model for an organisation's region; until then neither is made there, and rules are read by ZeeDo's model as described above. The same restriction applies to both: the provider keeps no copy and does not train on what it receives.

- **Reading rules.** When a person types a message that sets or changes a rule for a bot (with words such as "never", "always" or "from now on"), the typed text of that message and the words of the rules it concerns are sent, to tell whether it is a standing rule; so is the start of a request to a bot whose rules do not all fit in front of it. For this reading, what a person pastes, and anything a bot read, are not sent.
- **Checking a press.** Before a bot presses a button or another control that ZeeDo's own signals show could pay, order, sign, file, send, publish, delete or change an account, the control is checked, to tell what pressing it would do, so that ZeeDo can stop the press or ask a person first. Only labels are sent: the control's label, the labels of its form's fields and of a dialog it answers, the currencies shown, the site's domain and the pattern of the form's address (its path, without the query: a part written only in lower-case letters a to z, hyphens and underscores is kept as written, so it can be a name, such as a user name or a page's name, and every other part is replaced by a placeholder). Field values, the page's text and its links are never sent. Inside a label, every number written in digits (a house number, a postcode, an

amount, a date or a reference alike), email addresses, phone numbers, account and card numbers and web addresses written with their scheme (such as https://) or starting with www. are replaced by a placeholder before anything is sent. The other words of a label are sent as written: they may name people, for example the person a message is addressed to, or a street or a place, and an amount written out in words is sent as written. Most presses are not checked by a model at all. Pages of the banking, tax, health and government websites the site registry marks as sensitive, of a site a workspace has marked as sensitive, and of webmail are never sent to the United States: for an organisation of our EU region they are checked in the EU by ZeeDo's model, or not by a model at all. Only the answer and the site's full host name (such as shop.example.com) are kept, never the labels: for at most 30 days, in the organisation's region. Where a press cannot be checked, the bot does not make it on its own: a person presses the control or approves the press.

These two uses are not made for a workspace that uses the customer's own key. For an organisation of our EU region, an Owner or Admin can also turn on **Keep every model check in the EU** in Settings: then neither use goes to the United States. Rules are read by ZeeDo's model in the EU as described above, and presses are checked by ZeeDo's model in the same way or, until that check is available, only by ZeeDo's own word lists and the signals in the page. With this setting on, these two checks are made by ZeeDo's model in the EU even for a workspace that uses the customer's own key, and never by that key's provider.

Separately from the model a bot uses, we turn the summaries a bot keeps of what it did, the lines kept from its chat, its procedure documents and the text of each memory search into vectors (embeddings), so that memory can be searched by meaning, and we order what a search finds with a second model of our own. Both models run on our own servers in France, for every customer and whatever key it uses, and keep no copy of the text: a vector is stored with its memory item and deleted with it. Tool search: the names and descriptions of the tools a workspace has connected (which may be the vendor's own text) and the words a bot searches its tools with are turned into vectors by ZeeDo's own embedding model, the one that embeds memory text by default, on ZeeDo's servers; the vectors are kept with the workspace's tool registry and deleted with it. Where an Owner or Admin of a workspace chose an embedding model of the customer's own OpenAI account instead, that workspace's memory text goes to OpenAI under the customer's own contract, and OpenAI acts as the customer's processor (section 7.2).

5.3 **Long-lived memory.** Bots keep memory so they can work over months or years. **The customer controls how long memory is kept** and can delete it at any time. Our defaults and limits:

Layer	What it holds	How long it is kept
Working context	The current run's messages	For the run, then summarised into episodic memory. Checkpoints kept 30 days
Episodic memory (raw episodes)	A record of each step with a pointer to its full text	**At most 90 days**

Layer	What it holds	How long it is kept
Episode summaries	Short summaries of past steps	Until the customer deletes them or the workspace is deleted
Semantic graph	Facts about entities and relationships, each with a validity period	Until the customer deletes them or the workspace is deleted. When a fact stops being true it is marked invalid and kept as history; an erasure request deletes it (9.2)
Procedural memory	How-to documents, the bot's brief, recorded step-by-step flows	Until the customer deletes them or the workspace is deleted
Notes	Shared notes written by bots and users	Notes bots write 30 days by default; notes people write until deleted

Memory never crosses from one workspace to another. An erasure request finds a person's data in memory by the identifiers given, such as a name or an email address, wherever they appear in its text (9.2). Recorded flows keep secret field values as references to the vault, never as text.

5.4 **Screen recordings.** While a bot works on its screen, we save a still image of the screen every 2 seconds and one after each action the bot takes on it. In the bot's browser, the fields a page marks as password or payment card fields, payment forms embedded from payment providers, and fields the vault fills in are masked before any image is saved or shown, and an image whose masks cannot be worked out in time is masked whole. While a person takes over or teaches a bot, every input field of the page is masked as well. A field the page does not mark in one of these ways, and what is shown outside the bot's browser, is not masked. While a person completes a step a bot handed to them, recording stops. No keystrokes or clipboard content are recorded. Watching a bot's screen live is not recorded as a separate stream.

5.5 **Sensitive websites.** For pages of the banking, tax, health and government websites the site registry marks as sensitive, of a site a workspace has marked as sensitive, and of a site a person signed in to for a bot, the trace of each browser step keeps only the page's address, the step and the outcome. The page's text is left out of the trace, including the trace of the model request that read it. This applies to traces only: the run's checkpoints and its screen recordings keep what the bot saw for their own periods (section 6).

5.6 **Credentials.** Credentials in the vault are encrypted with keys held in AWS Key Management Service in Frankfurt. Bots use credentials through the tools a customer allows but cannot read them, and no tool returns a credential to a model. Only an Owner or Admin, signed in and confirming their sign-in again, can reveal a credential to themselves, and each reveal is recorded in the audit log.

5.7 **Webhooks.** Webhook deliveries to a customer's systems carry identifiers, the event type, times and the event's non-personal fields only.

5.8 **No training.** We do not use Customer Content to train or fine-tune any AI model, and ZeeDo's model uses only providers that do not train on it.

5.9 **Operational logs of bot machines.** For security and operation, the servers that run bot machines log what each machine did: the commands it ran, with their arguments, the paths of the files it wrote (not their contents) and the websites it reached, with the identifiers of the run. These logs are kept for at most 30 days (section 6) and are held with our infrastructure logs (7.6).

5.10 **Look-outs and checks.** Where a person turns on "Look out for me" for a bot, the bot reads, once a day, the mail and calendar apps the customer connected and allowed it to read, and may also read again later to check on something it did. These runs can only read: they cannot send, change or delete anything. What the bot reports is kept as a short note in the bot's chat: a title, a few sentences and references to where it read them (such as a message's or an event's identifier), never the text of a mail. A note is deleted after 30 days, or earlier when the person dismisses it; a person can switch look-outs off at any time.

6. How long we keep data

Each period applies on every plan, Free included, unless the row says otherwise.

Data	Retention
Run history (the step record, not files)	Per plan: Free 7 days, Starter 30, Pro 90, Business 365, Enterprise as agreed. A paid add-on extends it in 90-day steps
Conversations with bots	On every plan, until the bot or the workspace is deleted. An erasure of a person deletes the lines that name them
Tasks on the board that are done or cancelled	The run-history period of the plan. A task still open stays
Decisions and standing instructions on the board	Until a person deletes them or the workspace is deleted
Screen recordings	14 days on every plan. Business may extend this to 90 days and Enterprise to 400, never shorten it
Model context: run checkpoints and trace bodies	30 days, on every plan
Workflow history of finished jobs	At most 30 days after the job ends
Trace metadata (identifiers, timings, counts; no content)	13 months
Memory: raw episodes	At most 90 days
Memory: summaries, facts and relations, procedures, notes people write	Until the customer deletes them or the workspace is deleted
Memory: notes bots write	30 days by default
Files bots made	Until a person deletes them or the workspace is deleted; files that are screen recordings follow the recordings' 14 days

Data	Retention
Files a person hands a bot	A file that was never sent with a message or card is deleted at the first nightly pass at least a day after it was stored, so within two days. A file that was sent is kept as long as the message or card it was sent with, and deleted with it
Vault credentials and browser sessions	Until the customer revokes or deletes them; then deleted at once
Workspace machine snapshots	For the life of the workspace
Operational logs of the servers that run bot machines	At most 30 days
Export archives	7 days after they are written
Failed webhook deliveries	Retried for 24 hours, then kept 7 days for replay
API idempotency records	24 hours
Cost records, per-call detail	90 days. A run's total cost is kept as long as the run
Hourly and daily usage roll-ups	13 months
Daily and monthly totals of the usage of ZeeDo's model	25 months
Audit log and consent records, including approval decisions and standing approvals	400 days on every plan, and at most a month longer, because the log is deleted one calendar month at a time. The plan limits how far back the log can be searched, not how long rows are kept
The audit log's hourly write-once copy	400 days
Data of a removed Member	Deleted 30 days after removal. Audit rows keep only the pseudonymous member id
Expired sign-in sessions	30 days after expiry
An organisation after end of service	The clock in the Terms of Service, section 18: read-only with export days 0 to 60, deleted on day 90, purged from backups within 30 days after
Backups	Rolling, at most 30 days
Identifiers in an erasure or export request	Until the job completes
Invoices and credit notes	8 years from the end of the calendar year of issue, in write-once storage (§ 147 AO, § 257 HGB)

Data	Retention
Copies of the emails we send a customer about its contract, among them contract confirmations, notices of changes to the Terms, the AUP or the DPA, billing emails and notices of restrictions	6 years from the end of the calendar year in which they were sent (§ 147 AO, § 257 HGB)
Records of who accepted which version of the Terms, the AUP and the DPA, and of objections	For the life of the contract, then until the end of the third calendar year after it ended, and at least as long as the confirmation
Billing events received from our payment provider	400 days. The content of each event is removed after 30 days
Records of billing changes and of the billing emails we send (their kind, recipients and time)	400 days
Website and server logs	14 days
Infrastructure metrics	At most 13 months
Traces of our own services' requests (identifiers and timings)	At most 30 days
Application error reports	At most 90 days
Status-page subscriptions	Until you unsubscribe
Security disclosure reports	400 days
Support correspondence	24 months after the last message. A message that is a business letter, such as a cancellation, a switching request or a dispute about an invoice, and our answer to it: 6 years from the end of the calendar year in which it was received or sent (§ 147 AO, § 257 HGB)
Abuse and suspension records	400 days
One-per-person record of free AI usage	Until 400 days after the organisation is deleted; the network code 2 days after its day. An erasure request does not remove it earlier

7. Who receives data

7.1 **Sub-processors.** We use the service providers listed in Annex 3 of our DPA, which is also published on its own page at zeedo.cloud/legal/subprocessors. The list names each provider and each hosting site, with its location. We give at least 30 days' notice, by email to each customer's Owners and on the published list, before adding or replacing a sub-processor, including a model provider named on the list, except where a sub-processor must be replaced at shorter notice to keep the service secure or available (section 5.3 of the DPA). The messages sent to and from our mailboxes are held by our email provider.

7.2 **Model providers.** With the customer's own key (BYOK), the provider the customer chose receives prompts under the customer's own contract, as the customer's processor. With ZeeDo's model, which is the default, OpenRouter and the model provider named for each model on the sub-processor list receive them as our sub-processors. Every request is restricted to providers that keep no copy of prompts or responses and do not train on them, and for an organisation of our EU region requests are processed in the EU, except the reading of the rules people set for a bot and the check of what a control a bot is about to press would do, which may be served from the United States unless the organisation keeps every model check in the EU (5.2). Embeddings, and the ordering of memory search results, are computed by our own models on our own servers, and no model provider receives that text (5.2), unless an Owner or Admin of a workspace chose an embedding model of the customer's own OpenAI account: OpenAI then receives that workspace's memory text and searches under the customer's own contract, as the customer's processor. That text can contain personal data.

7.3 **Third parties a bot contacts.** A bot sends data where its customer instructs it to (for example, an email to a supplier). That is the customer's disclosure, not ours.

7.4 **Authorities and advisers.** We disclose data to authorities and courts where the law requires it, and to our tax advisers, lawyers and auditors, who are bound to confidentiality, where they need it to advise us.

7.5 We do not sell personal data.

7.6 **Traces and monitoring.** Langfuse (Langfuse GmbH, part of ClickHouse) holds the traces of bot runs that we export to it, in its EU region in Ireland: their bodies for at most 30 days and their metadata for at most 13 months (section 6). Langfuse's support may access them from the United States (8.2). Grafana Labs holds our infrastructure logs, metrics and alerts, and the traces of our own services' requests, in Grafana Cloud, in an EU region in Germany, including the operational logs of the servers that run bot machines (5.9).

8. Location and international transfers

8.1 **Where data is.** Customer data is stored and processed in the EU (France, Germany and Ireland). Bot machines, ZeeDo's own embedding and reranking models, the control-plane server and the primary file storage run at Scaleway in its Paris region (fr-par), France; a copy of the file storage is kept at OVH in Limburg, Germany; the database (Neon), the workflow service (Temporal Cloud), key management (AWS KMS) and transactional email (Amazon SES) run in Amazon Web Services' Frankfurt region (eu-central-1), Germany; the web application's functions run on Vercel in Frankfurt (fra1), Germany; and traces are kept by Langfuse in its EU region, in Ireland. Our infrastructure logs, metrics and alerts are kept by Grafana Labs in Grafana Cloud's EU region in Germany (7.6).

The exceptions are: the decision model that reads the rules people set for a bot and checks the labels of a control a bot is about to press, which may be served from the United States unless an organisation of our EU region keeps every model check in the EU (5.2); requests on ZeeDo's model of an organisation outside our EU region; the delivery of the web application's pages and responses through Vercel's global edge network, which passes them through the edge location

nearest the user; a model provider, or an embedding model of its own OpenAI account, that a customer chose with its own key, which processes data where that provider does; the places a customer's bots send data to on its instructions; and the recipients section 8.2 names, which process some data in a country outside the European Economic Area or may access it from there. Members' sign-in data is processed by WorkOS in the United States (8.2).

8.2 **Transfers.** Some recipients are in a country outside the European Economic Area (EEA), or may access data from one:

- **WorkOS,** our sign-in provider, processes Members' account and sign-in data in the United States.
- **Perplexity** runs the decision model that reads the rules people set for a bot and checks the labels of a control a bot is about to press, in the United States once ZeeDo uses it; it is reached through OpenRouter (5.2).
- **OpenRouter** is a United States company. For an organisation of our EU region it processes requests on ZeeDo's model at its EU endpoint, in the EU; for an organisation of another region, requests may be processed outside the EU.
- **Amazon Web Services, Microsoft (**which serves ZeeDo's default model at an EU endpoint), **Vercel, Sentry, Neon, Temporal, Grafana Labs** and **Stripe** process data for us in the EU but belong to groups based in the United States, from where they may be able to access it. Stripe may also process billing and payment data in the United States.
- **Langfuse**'s support may access traces from the United States.
- **Better Stack,** which runs our uptime checks and our status page, processes the email addresses of status-page subscribers in the United States and other countries.

For every one of these recipients except Better Stack we rely on the European Commission's Standard Contractual Clauses (Implementing Decision (EU) 2021/914), in the recipient's data processing terms or, for the model providers OpenRouter engages, in OpenRouter's. In addition, Amazon Web Services, Microsoft, Stripe, Vercel, Sentry, Grafana Labs and Neon state that they are certified under the EU–US Data Privacy Framework, for which the European Commission adopted an adequacy decision on 10 July 2023. For WorkOS, OpenRouter, Perplexity and Temporal we rely on the Standard Contractual Clauses alone. Better Stack's data processing terms rely on its certification under the EU–US Data Privacy Framework, and on the Standard Contractual Clauses where the Framework does not apply. Any other recipient on the sub-processor list that processes data outside the EEA, or can access it from there, does so under the Standard Contractual Clauses unless the country has an adequacy decision. Where a customer uses its own key or its own OpenAI account, the transfer to that provider is the customer's, under its own contract. You can ask for a copy of our safeguards at privacy@zeedo.cloud.

9. Your rights

9.1 You have the right of access (Article 15 GDPR), to rectification (Article 16), to erasure (Article 17), to restriction of processing (Article 18) and to data portability (Article 20). We do not ask for your consent for any processing in section 4; if we ever do, you may withdraw it at any time,

with effect for the future.

Your right to object (Article 21 GDPR). Where we process your data on the basis of legitimate interests (Article 6 (1) (f) GDPR), you may object at any time, on grounds relating to your particular situation. We then stop that processing, unless we demonstrate compelling legitimate grounds that override your interests, rights and freedoms, or need the data to establish, exercise or defend legal claims.

9.2 **Erasure.** A customer files an erasure request in the product for a person whose data its bots process; if you write to us about such data, we pass your request to the customer (3.3). An erasure request runs one job across the stores that hold the text of bot memory, chats, runs, cards and approvals, automations, standing approvals and connections, traces, recordings, checkpoints and the files people hand bots. Within 72 hours it permanently deletes what the identifiers given find: in bot memory the matching raw episodes, notes, facts and relations, and the summaries that name the person; the chat lines that name them, each replaced by a marker that a line was erased; the text of the runs, cards and approvals that name them, and those runs' checkpoints and screen recordings; trace bodies; the files handed in those messages and runs; and the text of automations, standing approvals' conditions and connection account names. It then produces a report. Data the identifiers cannot find, such as a face or a name visible only in a screen recording or in a checkpoint of a run whose text never names the person, is deleted when its own period ends (14 days for recordings, 30 days for checkpoints). Deleted data leaves backups within 30 days. Audit and consent records are kept and listed in the report as retained with a reason.

The job does not search the files bots made, the disks of bot machines and their snapshots, the operational logs of bot machines (5.9), the history of finished workflows, or what people wrote to set a workspace up, such as a bot's name or goal or a connection's label. The customer can delete the files and change or delete the settings itself; the logs and the workflow history are deleted when their periods end, and a machine's disk and its snapshots are kept for the life of the workspace (section 6). The full text of each raw memory episode is stored sealed and cannot be searched for a name, so the job deletes the full texts of every raw episode in every workspace of the organisation, and keeps the short summaries that do not name the person. The report lists each kind of data the job keeps, with its reason.

We keep audit and consent records for 400 days because they may be needed to establish, exercise or defend legal claims (Article 17 (3) (e) GDPR). Audit records identify the person who acted by a pseudonymous member id, and the name is looked up only when a record is displayed, so once a Member's account data is deleted their name no longer appears beside their audit records.

An erasure request does not remove the one-per-person record of free AI usage early (section 6). It holds only one-way codes, it is still needed for its purpose, which is to stop the free amount from being claimed again by deleting and re-creating an account, and our compelling legitimate grounds for keeping it override your interests (Article 17 (1) (a) and (c) and Article 21 (1) GDPR).

For data we control (section 4), we answer an erasure request within the time 9.5 states and delete your data where Article 17 GDPR requires it. We keep what the law requires us to keep,

what we need to establish, exercise or defend legal claims, and the records this section names and the abuse and security records that are still needed for their purpose, each for its period in section 6. While you are a Member of an organisation, we need your account data to provide its service; it is deleted 30 days after you are removed from the organisation.

9.3 **Corrections to memory.** An ordinary correction to a bot's memory keeps the earlier fact as invalid history, so that the bot knows what was true when. An erasure request is different: it deletes (9.2).

9.4 **Export.** While an organisation is active, and while it is read-only until day 60 of the end-of-service clock, an Owner or Admin can export its data as one archive signed by ZeeDo. The archive is produced within 24 hours and can be downloaded for seven days. Section 18 of the Terms of Service lists what the archive contains. Files bots made can be downloaded one by one in the product; they are not yet in the archive. For the data you gave us that we process to perform a contract with you, write to privacy@zeedo.cloud and we send it to you in a structured, commonly used and machine-readable format (Article 20 GDPR).

9.5 **How to ask.** Write to privacy@zeedo.cloud. We answer within one month. Where a request is complex, or we receive many, we may extend this by up to two further months; we then tell you so within the first month, with the reason (Article 12 (3) GDPR). We may first ask you to confirm your identity. If your request concerns data a customer's bot processed, see 3.3.

9.6 **Complaints.** You may complain to a data protection supervisory authority, in particular in the EU member state where you live or work, or where you think an infringement took place. Our competent authority is named in 1.3.

10. Security

10.1 We protect data with the technical and organisational measures described in Annex 2 of our DPA. Each customer workspace runs in its own hardware-isolated microVM, not a shared container.

10.2 If a personal-data breach affects customer data, we notify the affected customer without undue delay, and where feasible within 48 hours of confirming it. Where a breach affects data we control, we notify our supervisory authority unless the breach is unlikely to result in a risk to you, and we tell you where it is likely to result in a high risk to you (Articles 33 and 34 GDPR).

11. Automated decisions

11.1 We take no decision based solely on automated processing that produces legal effects concerning you or similarly significantly affects you, with one exception. To stop abuse or a security threat, our systems may restrict a bot, a workspace or an organisation at once (Terms of Service, section 10.2). Such a restriction concerns the customer's account. Where it is nonetheless a decision about you within Article 22 GDPR, it is necessary to perform our contract with you where you are the customer (Article 22 (2) (a) GDPR), and it protects our other customers and third parties. You can have it reviewed by a person, give your view and contest it by appealing to abuse@zeedo.cloud; a person at ZeeDo reviews the appeal afresh, and where

another person at ZeeDo can review it, that person did not take the original decision. Our fair-use check only flags a workspace: a person at ZeeDo decides on any fair-use suspension, after a written notice and a 72-hour cure period.

11.2 Customers' bots may act on data about you under the customer's control; see 3.3.

12. Changes

We publish changes to this policy on this page with a new "Last updated" date, and we will keep earlier versions available at their own addresses. Before a material change takes effect, we will tell account holders by email and in the product. This policy describes our processing of personal data; it is not part of the contract with a customer, which the Terms of Service and the DPA govern.