

# Privacy Policy

This notice explains what happens to your personal data when you visit upreltic.com, create an account, and use Uprelic. It is written to be read: each processing activity below states **what we do with the data, on what legal basis, and how long we keep it**.

Uprelic is an AI product. That means some of what you type, upload or ask for is sent to the companies that operate the AI models we use. Section § 2.3 says exactly what leaves our servers and what does not. We think that is the most important part of this notice, so we have not buried it.

English is the drafting language of this notice; a German version is published alongside it. This notice is not part of any contract — it informs you (§ 18.1 of our Terms of Service).

---

## § 1 Controller and contact details

---

The controller for the processing described here, within the meaning of Art. 4 No. 7 GDPR, is:

**Uprelic GmbH**, represented by Marco Herzog Liebenwalder Str. 16 13347 Berlin Germany

Email: [info@upreltic.com](mailto:info@upreltic.com)

**For data-protection matters, including any of the rights in § 6, write to [privacy@upreltic.com](mailto:privacy@upreltic.com).** We are not required to appoint a data protection officer and have not appointed one, so that address reaches the people who actually operate the systems described below.

Where a company or other organisation provides your workspace, that organisation — not us — is the controller for the content in it. See § 2.17.

---

## § 2 What we process, why, and for how long

---

### 2.1 Visiting the site and using the app: server logs

**What.** When your browser or the app talks to our servers, the request is logged: IP address, date and time, the URL requested, HTTP status and response size, referrer, and the user-agent string your browser sends.

**Why.** To deliver the site, to keep it stable, and to detect and investigate attacks, abuse and faults.

**Legal basis.** Art. 6(1)(f) GDPR. Our legitimate interest is operating a service that works and is not overrun by automated abuse. Where the request is part of using your account, the processing is also necessary to perform our contract with you, Art. 6(1)(b) GDPR.

**Retention.** Log lines collected into our log store are deleted **90 days** after they are written. The raw log files the web server writes for itself are not yet covered by that schedule.

Our site is reached through a network and security provider that terminates TLS at the edge and filters malicious traffic before it reaches us. That provider therefore sees connection data, including your IP address. See § 3 and § 4.

## 2.2 Account and registration

**What.** To create an account we process your **email address**. You sign in with a one-time link sent to that address, or with **Google** or **Apple**. If you sign in with Google or Apple, we receive the identifiers those services return for the sign-in. We do not set or store a password. You may add a display name, a handle and a profile picture.

**Where we got it, if not from you.** Where you sign in with Google or Apple, we receive the sign-in data from that provider rather than from you directly (Art. 14 GDPR). The categories are the account identifier, the email address associated with it, the name on that account, and — from Google — a link to your profile picture. Data also reaches us from Slack rather than from you where you link a Slack identity to your account (§ 2.18).

**Why.** To create and operate your account, to authenticate you, and to provide the services you have asked for.

**Legal basis.** Art. 6(1)(b) GDPR — performance of the contract you conclude when you register.

**Whether you have to give us this.** An email address is required to conclude the contract: without one we cannot create an account or sign you in, and there is no other way for us to do it. Everything else — display name, handle, profile picture, billing name and address beyond what an invoice needs — is voluntary, and leaving it out costs you nothing but the feature it feeds.

**Retention.** For as long as your account exists. When you delete your account we act as described in § 6.5: the account record itself is kept but stripped of the data that identifies you, so that content other people still hold — a chat someone shared with you, an invoice we must keep — does not break or become unattributable.

## 2.3 AI processing: your prompts, files and context

This is the block that matters most, and it is the one many AI privacy notices leave out. We are the controller for it and we can describe it plainly.

**What leaves our servers.** When you send a message, run a skill or ask the assistant to do something, we transmit to the operator of the model you have selected: your message, the parts of the conversation and of your project that are needed to answer it, the text extracted from files you attached, and — where the model is a vision model — the images concerned. Where the

answer draws on your memory profile (§ 2.8), that profile text is part of the prompt too. Two background steps carry a little more: when we derive memory from a chat or a project, the display name on your account is part of that prompt, and where an automation runs an AI agent step (§ 2.5), the prompt also carries the name and the email address on your account so that the step knows on whose behalf it is acting.

**What does not leave our servers.** Your files themselves are stored on our own infrastructure. **Search indexing and embedding run on our own servers** for everything you send us — the model that turns your content into vectors, and the model that re-ranks search results, run in our own process on our own hardware. That holds for the batch case as well: when we change the embedding model and have to re-index the whole corpus, the job runs on our own GPU, like every other embedding run. Code the assistant runs for you runs **on our own servers**.

**Which providers.** Which company receives a prompt depends on the model in use. The current model providers are named, by name and by role, on our [Subprocessors & AI Providers page](#). That page is kept current; naming providers there rather than here is what lets us keep it accurate. Interactive chat runs by default on a model operated in the United States; you can select another model in the model picker, and further processing steps — planning, summarising, generating a chat title, extracting memory, optical character recognition on scanned documents, image generation — may run on a different provider from that page.

**No training on your content.** We do not use your content to train or fine-tune AI models, and the providers we engage do not train on it either (§ 14.3 of the Terms). Where a provider offers that as a contract term we take it; otherwise we rely on the provider's standard commitment for business API use. Which of the two applies is stated per provider on the Subprocessors page. Either way it is a commitment we hold them to, not something we can technically verify inside their systems.

**Legal basis.** Art. 6(1)(b) GDPR — this processing *is* the service you contracted for. Where you have the assistant process personal data about other people, please read § 2.17 and § 13.4 of the Terms: you decide what to submit, and you need the rights to submit it.

**Retention.** Your chats, files and projects stay in your workspace until you delete them or delete your account. Prompts and model responses are also written to our own tracing system — see § 2.11.

## 2.4 Web search

**What.** When the assistant needs current information, it formulates a search query and sends **that query, and only that query**, to a web-search provider. The query is generated from your request, so it can contain what you asked about, including a person's name. It does not carry your name, your account, or any identifier that would let the provider connect it to you. The provider returns a list of results; the pages themselves are then fetched by **our own browser infrastructure**, not by the search provider on our behalf.

**When the page is a video.** If the assistant opens a video page and the video has no captions to read, we download its audio track and send it to a speech-to-text provider named on the Subprocessors page so that the assistant can use what was said. That audio belongs to the video rather than to you, but it is a transfer to a third party and so it belongs in this notice.

**Why.** To answer questions that require information from the public web.

**Legal basis.** Art. 6(1)(b) GDPR — performance of the contract.

**Retention.** On our side, the query itself is retained only as part of the trace described in § 2.11. Each search provider keeps its own record of the queries we send it; for how long is stated per provider on the Subprocessors page.

## 2.5 Automations, transcription and image analysis

**What.** An automation is a sequence of steps you configure that runs without you being present. Depending on the steps you build, an automation may **transcribe audio**, **analyse images**, or run text steps such as summarising or translating. To do that, the audio, the image, or the text concerned is transmitted to the model provider serving that step.

**Why.** To run the automation you have configured.

**Legal basis.** Art. 6(1)(b) GDPR.

**Retention.** The input and the result are stored with the automation run in your workspace until you delete the run or the automation. The provider's own retention is governed by our contract with it; see § 3 and § 4.

**A caution that belongs here.** If you have an automation transcribe a recording or analyse a photograph, you are responsible for having the consent of the people recorded or depicted (§ 9.5 of the Terms). In Germany, recording non-public speech without consent is a criminal offence (§ 201 StGB).

## 2.6 Connected Accounts and integrations

**What.** You can connect third-party accounts — a mailbox, a calendar, a project tracker, a file store. The connection is made through an **integration broker**, which is named on the Subprocessors & AI Providers page.

**Where the credentials live.** For accounts you connect through the integration catalogue, the OAuth access tokens are **held by that broker, not by us**. We store only a reference to the connection, which service it is, and whether it is active.

**Two exceptions.** Where you install our **Slack** app into a workspace, that workspace token is held **by us** rather than by the broker, because the app is ours; it is a surface of its own and what it processes is set out in § 2.18. And for a small number of integrations we operate directly, you enter the credentials yourself — an API key or a client secret — and we store them tied to your account. We hold no passwords for any third-party account.

**What we fetch.** Data is fetched from a connected service **when you instruct it** — when you ask the assistant to look something up, or when an automation you built runs a step against that service. Read access is implicit in the connection; **write and delete actions require a separate, explicit grant from you**, which you can give for a single conversation or until you revoke it.

**Google data.** Where you connect Google services, our use and transfer of data received from Google APIs follows the Google API Services User Data Policy, including its Limited Use requirements.

**Legal basis.** Art. 6(1)(b) GDPR — performance of the contract; the connection itself is made on your instruction.

**Retention.** Until you disconnect the integration or delete your account. Disconnecting revokes our access at the broker and, where you entered the credentials yourself, stops us using them; those credentials go when you delete your account, or sooner if you ask us at [privacy@uprelic.com](mailto:privacy@uprelic.com). Data already brought into your workspace stays there until you delete it (§ 9.6 of the Terms).

## 2.7 Dictation

**What.** If you dictate instead of typing, your microphone audio is streamed to our servers and transcribed there.

**Where. On our own hardware.** The speech-recognition model runs in our own process on our own machines. Your voice is not sent to a third-party transcription service.

**Why.** To turn speech into the text you then send.

**Legal basis.** Art. 6(1)(b) GDPR.

**Retention. None.** The audio is held in memory only for as long as it takes to produce the transcript, and is not stored afterwards. Only the resulting text is kept, as part of your message.

## 2.8 Memory and personalisation

**What.** So that you do not have to repeat yourself, we derive short summaries from your conversations — what you are working on, how you like answers formatted, facts you have told the assistant — and store them with your account, with the individual chat and — where you work in a project — with that project; we also keep a short work profile (your role, seniority and company context) once you have confirmed it in a guided chat. These summaries are generated from **your conversations**, which includes material that reached the conversation from a file you attached or from a chat someone shared with you.

**Why.** To personalise your workspace and make later answers more useful.

**Legal basis.** Art. 6(1)(f) GDPR. Our legitimate interest, and yours, is an assistant that does not start from zero every time. You can object to this processing at any time under § 6.4.

**Retention and control.** Deleting a chat moves it to Trash rather than erasing it, so memory derived from that chat continues to exist for as long as the chat does. Memory held with your account is deleted when you delete your account; memory attached to a chat or a project that stays in existence because you shared it stays with it, unless you ask during deletion for the shared items to go as well (§ 6.5). What the data export does and does not carry is set out in § 6.6.

## 2.9 Payments and billing

**What.** Paid plans and top-ups are processed by **Stripe**. Card details are entered directly into Stripe's own payment fields; **we never see or store your card number**. We transmit to Stripe your email address, an internal reference to your account, and — where you enter them — your billing name, billing address and VAT identification number, which are needed to issue a correct invoice and to calculate tax. Stripe returns to us the status of the payment, the subscription, and the links to your invoices.

**Stripe's own role.** For fraud prevention and its own regulatory duties, Stripe acts as an independent controller and processes payment data under its own privacy policy.

**Why.** To conclude and perform the contract, to take payment, and to meet our tax and commercial-law duties.

**Legal basis.** Art. 6(1)(b) GDPR for the payment itself; Art. 6(1)(c) GDPR for the retention of accounting records.

**Retention.** Invoices and the underlying accounting records are kept for **eight years**, as required by § 147 AO, § 257 HGB and § 14b UStG. Subscription and credit records are kept for as long as they are needed to account for your plan, and the tax-relevant parts fall under the same eight-year rule. Billing records are deliberately **not** deleted when you delete your account — we are not permitted to delete them.

## 2.10 Emails we send you

**What and why.** We send email that the service itself requires: sign-in links, confirmation of your order (§ 312f BGB), confirmation of a cancellation or a withdrawal (§ 356a BGB), notice that your data export is ready, notice that your account has been deleted, invitations someone has sent you, and notifications about activity that concerns you — a mention, a share, a request for access.

**Legal basis.** Art. 6(1)(b) GDPR for contractual and statutory mail; Art. 6(1)(f) GDPR for activity notifications, our legitimate interest being to tell you about things in your workspace that need you.

**Product information.** Where we email you about changes to the service you already use, we do so on the basis of Art. 6(1)(f) GDPR, our legitimate interest being to keep existing customers informed about the service they rely on.

Where you have bought a paid plan or a top-up, we may also email you about our own similar services, on the basis of Art. 6(1)(f) GDPR in conjunction with **§ 7(3) UWG**; otherwise we send such mail only if you have asked for it. **You can object to this use of your address at any time, at no cost beyond the transmission costs at basic rates**, by writing to [privacy@uprelic.com](mailto:privacy@uprelic.com).

**We do not run a newsletter, and we do not track our emails.** There are no tracking pixels and no click tracking in the mail we send.

**Transport.** Mail is sent through **Amazon SES in the eu-central-1 (Frankfurt) region**.

**Retention.** Delivery logs are kept as part of the operational logs described in § 2.1. The content of a message is whatever we sent you.

## 2.11 Error monitoring and AI tracing — on our own servers

**We use no product-analytics service. How you use the product is not tracked for analytics.** There are no page-view or interaction events, no session recording, and nothing that measures which features you use or where you get stuck — not on a third-party platform, and not on one of our own either. What is left observes how the product behaves rather than how you use it, and it runs on our own infrastructure rather than someone else's: error monitoring and AI tracing are both first-party.

**Error monitoring.** When something breaks, an error report is written to our self-hosted error-monitoring system: the fault and its stack trace, the release, and **account data that identifies you**, so that we can help when you tell us something failed. Recent log lines are attached to the report as context, so whatever those lines contain travels with it. There is no session replay.

*Legal basis:* Art. 6(1)(f) GDPR — our legitimate interest in finding and fixing faults in a service people rely on. *Retention:* for as long as the record remains useful for diagnosing faults; we have not set a fixed deletion period.

**AI tracing.** So that we can debug wrong or failed answers and investigate abuse, we record **the prompts sent to models and the responses received**, linked to your account and to the conversation, in our own tracing system on our own infrastructure. **These traces are not used to train or fine-tune any model** (§ 18.2 of the Terms). Deleting an individual chat removes it from your account, but not from these internal records.

*Legal basis:* Art. 6(1)(f) GDPR — our legitimate interest in operating an AI service safely and correcting it when it misbehaves. *Retention:* **for as long as your account exists**. When you delete your account these records are deleted with it; anonymous technical measurements — how many tokens, how long a request took, what it cost — that no longer identify you may be kept.

## 2.12 Cookies and device storage

**Strictly necessary.** We set what the service cannot work without: the sign-in cookie that keeps you signed in, a cookie storing your language choice, a cookie holding the country your request came from so that prices are shown in the right currency, and a short-lived cookie carrying a one-time value that protects the Slack connection flow against request forgery. In your browser's local storage we keep interface preferences — sidebar layout, chosen model, view modes — short-lived caches, and drafts of forms you have begun filling in so that you do not lose them, which can include what you have typed about a person you are inviting.

*Legal basis:* § 25(2) No. 2 TDDDG — these are strictly necessary to provide the service you requested — together with Art. 6(1)(f) GDPR for the associated processing of personal data, our legitimate interest being to keep you signed in and to show the service in the language and currency you expect.

**Nothing optional, and therefore nothing to consent to.** We set no optional cookies and load no optional scripts. § 25(1) TDDDG requires consent for storage on your device that is not strictly necessary; we have none, so there is nothing here for you to agree to and **no consent banner to click away**. If that ever changes we will ask you first, and this section will say what is being asked before anything is set.

**Retention.** The sign-in cookie lasts **up to 400 days**, or until you sign out; the language cookie one year; the country cookie 30 days; the Slack flow cookie a few minutes.

## 2.13 Support and contact

**What.** If you write to [info@uprelic.com](mailto:info@uprelic.com) or [privacy@uprelic.com](mailto:privacy@uprelic.com), we process your message, your address and anything you choose to put in it, in order to answer you. We operate no support chat widget and no third-party helpdesk — mail comes to us directly.

**Legal basis.** Art. 6(1)(b) GDPR where your enquiry concerns your contract; otherwise Art. 6(1)(f) GDPR, our legitimate interest in answering people who write to us.

**Retention.** Correspondence is kept for as long as needed to deal with the matter and to show that we did. Where a message forms part of a commercial transaction it falls under the statutory retention rules in § 2.9.

## 2.14 Forums and public content

**What.** Uprelic has project forums. **A post or comment you make in a public forum is public:** it can be read by people who are not signed in, it is served in the page source for search engines, and it can appear in search results. Your display name, your handle and your profile picture are shown with it.

**Why.** To operate a public discussion space.

**Legal basis.** Art. 6(1)(b) GDPR — publishing what you have chosen to publish is part of the service; and Art. 6(1)(f) GDPR for making public content findable.

**Retention.** Until you delete the post, or until the project it belongs to is deleted. If you delete your account, posts you made inside other people's projects remain, but they are no longer connected to an identifiable profile, because the profile they pointed to has been stripped of identifying data (§ 6.5). If you want a specific public post removed, tell us and we will remove it.

## 2.15 Using Uprelic without an account

**What.** You can open a share link, read public content, and continue a shared conversation as a guest, without registering. In a guest session we process the content of that session and the technical data in § 2.1. No account is created, and we do not set a cookie to recognise you across sessions.

**Legal basis.** Art. 6(1)(b) GDPR — a guest session is itself a contract (§ 3.5 of the Terms).

**Retention.** A guest conversation is stored so that it can be continued.

## 2.16 Records we keep to prove what was agreed

**What.** Consumer law requires us to be able to show which terms you accepted and when, and to record cancellations and withdrawals. We therefore store: the document and version you accepted, with the time and the language; and, for a cancellation under § 312k BGB or a withdrawal under § 355 BGB — the latter submitted through the withdrawal function that § 356a BGB requires us to provide — the declaration you submitted together with **your IP address and browser user-agent**, which are what make the declaration evidentially useful. We identify the account from the address you name; a handle or billing postcode, if you give one, is checked only as corroboration and can never cause a declaration to be rejected. We record which account we matched, how well it was corroborated, and what we did about the contract.

**One optional question, which is not part of the declaration.** The cancellation form also asks why you are leaving, as a short list of fixed answers — too expensive, not using it enough, missing features, and so on. Answering is voluntary, it is stored with the record, and it has no bearing whatsoever on the cancellation, which takes effect either way. We use it only to understand why people leave. **Legal basis:** Art. 6(1)(f) GDPR, our legitimate interest in learning where the product falls short. Leave it blank and nothing follows.

**Why.** To comply with our statutory duties, and to be able to prove compliance if it is disputed.

**Legal basis.** Art. 6(1)(c) GDPR — compliance with a legal obligation — and Art. 6(1)(f) GDPR, our legitimate interest in being able to establish and defend legal claims.

**Retention.** Until the end of the third calendar year following the declaration, which is the ordinary limitation period under §§ 195, 199 BGB, and longer where the record forms part of accounting documents subject to the eight-year rule in § 2.9.

## 2.17 Where a company provides your workspace

If your workspace is provided by a company, an employer or another organisation, then for the content in that workspace **that organisation is the controller and we act as its processor**, on its instructions and under a data processing agreement (Art. 28 GDPR). In that case, questions about that content — access, correction, deletion — are for that organisation to answer, and its own privacy notice applies to it. Ask them, not us; we will refer you to them.

We remain the controller in our own right for the things that are ours to decide: operating and securing the platform, our own billing relationship, statutory records, and aggregated technical measurements that do not identify individual users. Those are described in this notice.

## 2.18 The Uprelic Slack app

**What.** You can install our Slack app into a workspace and talk to the assistant from Slack. Several distinct things follow from that.

**The installation itself.** When the app is installed we store the workspace's installation record — which workspace it is, the credentials Slack issues for it, and the permissions granted — together with the identity of the person who installed it. Most of that is data about the workspace rather than about you, but whoever performed the installation is identifiable from it.

**The link between your two identities.** So that something you say in Slack can act on your Uprelic account, we store the link between your Slack identity and your Uprelic account, with a snapshot of the name and the email address on your Slack profile, and how and when the link was made. A link is made either because you confirmed it in the web app while signed in, or — where your workspace lets the app read verified profile emails, and yours matches exactly one Uprelic account — **automatically, without your being asked**.

**Where those two came from.** The name and the email address come from **Slack, not from you**, which makes them Art. 14 GDPR data. The source is your Slack workspace's own user directory; the categories are your Slack display name and your Slack profile email. We hold them so the app can show you which identity is linked, and so that there is a way to reach you if your Uprelic address stops working.

**What you say to it.** A thread you open with the assistant in Slack is mapped to an Uprelic chat. **The conversation itself is an ordinary Uprelic chat and is covered by § 2.3** — it is not kept a second time in a Slack table.

**Why.** To provide the Slack app that was installed and that you chose to use.

**Legal basis.** Art. 6(1)(b) GDPR — performance of the contract — for the installation, for a link you confirm yourself, and for the conversation. For the **automatic** link from a verified Slack email address, Art. 6(1)(f) GDPR: our legitimate interest, and yours, in a message sent from Slack reaching the right account instead of no account at all. You can object to that under § 6.4, and you can remove a link at any time.

**Retention.** A link lasts until you remove it. The installation record lasts until the app is removed from the workspace: at that point we mark it revoked and stop calling that workspace, though the record itself stays behind, holding credentials Slack has already invalidated. Thread mappings go when the chat or the account they point at goes. Deleting your account should end all of this — § 6.5 records honestly where it currently does not.

---

### § 3 Who receives your data

---

Inside our company, only those people who need it to do their job have access.

Outside it, we pass data to the following categories of recipient:

- **Hosting and infrastructure** — the servers on which your content is stored at rest.
- **Network and security** — the provider that fronts our site, terminates TLS and filters malicious traffic.
- **AI model providers** — the companies that operate the models processing your prompts and files (§ 2.3).
- **Web search** — the provider that answers search queries (§ 2.4).
- **Speech-to-text and image analysis** — the providers serving those steps in automations, and the transcription of video audio described in § 2.4.
- **Integration broker** — the service that holds the OAuth connections to your third-party accounts (§ 2.6).
- **Payments** — Stripe (§ 2.9).
- **Email delivery** — Amazon SES, eu-central-1 (§ 2.10).
- **App distribution** — Apple and Google, where you install the app from their stores.
- **Advisers and authorities** — our tax advisers and auditors, and public authorities where we are legally required to disclose.

**The individual companies behind these categories, what each one does, where it sits and on what legal footing, are listed on the [Subprocessors & AI Providers page](#).** We keep that page current, which is why the names are there rather than here. The page is informational; it does not form part of any contract.

One category is an exception we would rather flag than have you discover: the providers serving speech-to-text and image analysis are not currently named on that page because we are in the process of replacing them. Write to [privacy@uprelc.com](mailto:privacy@uprelc.com) if you want to know who they are today.

Systems that a company of our size would usually buy as a hosted service — our database and authentication, our vector search, our error monitoring, our AI tracing, our browser rendering and our code sandboxes — **run on our own servers**. They are therefore not third-party recipients at

all. The one such system we would otherwise have bought, product analytics, we simply do not run.

---

## § 4 Transfers to countries outside the EU/EEA

---

Some of the recipients in § 3 — in particular the AI model providers — are established in the **United States**. Where we transfer personal data to a country outside the EU/EEA, **we do so only where at least one of the following applies**:

1. the European Commission has decided that the country ensures an adequate level of protection (Art. 45 GDPR);
2. the recipient is certified under the **EU–US Data Privacy Framework**; or
3. we have concluded the European Commission's **standard contractual clauses** with the recipient (Art. 46(2)(c) GDPR), supplemented where necessary by additional measures.

You may request a copy of the safeguards in place by writing to [privacy@uprelc.com](mailto:privacy@uprelc.com).

| Processing                                                               | Where          | Transfer mechanism                                                              |
|--------------------------------------------------------------------------|----------------|---------------------------------------------------------------------------------|
| Hosting, database, search index, error monitoring, AI tracing, sandboxes | Germany        | No transfer — servers we rent at Hetzner in Germany and operate ourselves       |
| Email delivery                                                           | EU (Frankfurt) | No transfer                                                                     |
| Network and security edge                                                | USA / global   | Data Privacy Framework and standard contractual clauses                         |
| AI models for chat, documents and images                                 | USA            | Standard contractual clauses                                                    |
| Speech-to-text and image analysis, in automations and for video audio    | USA            | Standard contractual clauses                                                    |
| Web search                                                               | EU / USA       | Standard contractual clauses; the query carries no identifying data — see § 2.4 |
| Integration broker                                                       | EU / USA       | Standard contractual clauses                                                    |
| Payments                                                                 | USA / EU       | Standard contractual clauses; Stripe also acts as an independent controller     |
| App distribution                                                         | USA            | Data Privacy Framework                                                          |

---

## § 5 Automated decision-making

---

**There is one automated decision, and it runs whenever a sign-in link is requested — when you register and when you sign in again.** To keep automated sign-ups and throwaway

accounts off the platform, the request is refused automatically if the email domain appears on a published blocklist of disposable-email providers, if a hidden form field that a human never sees has been filled in, or if the form was submitted impossibly fast. Where one of these matches, no sign-in link is sent and, where the request was a registration, the account is simply not created.

That is a decision based solely on automated processing within the meaning of **Art. 22(1) GDPR**. It is permitted under **Art. 22(2)(a) GDPR** because it is necessary for entering into, or performance of, a contract with us: without it we could not offer the service on these terms. No profiling of your behaviour is involved, and no special categories of data are used.

**Legal basis for the checks themselves:** Art. 6(1)(f) GDPR — our legitimate interest in keeping automated sign-ups and throwaway accounts off the platform (Recital 47); Art. 22(2)(a) GDPR is what permits the decision to be taken automatically.

**How you will know, and what to do about it.** Of the three grounds, only the blocklist one tells you what happened: you get an error saying the address cannot be used. The other two do not. If the hidden field is filled in or the form is submitted faster than a person could manage, the page shows the same "check your email" confirmation as a successful sign-up and no link is ever sent. The same three checks run when you ask for a sign-in link for an account you already have, so a returning user can be caught by them too; and where a domain has been added to the blocklist since you registered, an existing account on that domain can no longer be reached by sign-in link at all. We would rather write that down than let you conclude your mail was lost.

**So: if you asked for a sign-in link and no email arrived, write to [info@uprelc.com](mailto:info@uprelc.com).** A person will look at it, and you can state your position and contest the decision — that is what Art. 22(3) GDPR requires, and we mean it literally. The usual cause is a legitimate address on an over-broad blocklist, or a form filled in faster than our threshold expects, and both are things we can fix.

**Beyond that, no.** We do not use automated decision-making or profiling that produces legal effects concerning you or similarly significantly affects you. Deciding which AI model answers a request, or personalising your workspace from memory (§ 2.8), is not a decision about *you* in that sense.

---

## § 6 Your rights

---

You have the following rights in respect of your personal data. **They are free of charge**, and exercising one is never held against you.

### 6.1 The rights themselves

- **Access** (Art. 15 GDPR) — confirmation of whether we process data about you, a copy of it, and the information in this notice as it applies to you.

- **Rectification** (Art. 16 GDPR) — correction of inaccurate data and completion of incomplete data.
- **Erasure** (Art. 17 GDPR) — deletion, where one of the grounds in Art. 17(1) applies and no retention duty stands in the way. Where we must keep something — invoices, for eight years — we will tell you what and why.
- **Restriction** (Art. 18 GDPR) — that we hold data without using it, for instance while a dispute about its accuracy is resolved.
- **Data portability** (Art. 20 GDPR) — the data you gave us, in a structured, commonly used, machine-readable format, and transmission to another controller where technically feasible.
- **Objection** (Art. 21 GDPR) — see § 6.4, which is set apart because the law requires it to be brought to your attention explicitly.
- **Withdrawal of consent** (Art. 7(3) GDPR) — see § 6.3.
- **Complaint to a supervisory authority** (Art. 77 GDPR) — see § 6.7.

## 6.2 How to exercise them, and how fast we answer

Write to [privacy@uprelc.com](mailto:privacy@uprelc.com). We answer **within one month** of receiving your request; if a request is complex, or if there are several, we may extend that by two further months and will tell you within the first month if we do (Art. 12(3) GDPR).

If we have reasonable doubts about who is making a request, we may ask for further information to establish your identity (Art. 12(6) GDPR). We ask for the least that will do the job — normally, that the request comes from the address the account is registered to.

## 6.3 Withdrawing consent

Where processing rests on your consent, you may withdraw it at any time with effect for the future, and withdrawal does not affect the lawfulness of what was done before it. **As things stand, none of the processing described in this notice relies on your consent.** We run no analytics and set no optional cookies, and everything above rests on the contract, on a legal obligation, or on our legitimate interest. If we introduce something that needs your consent, we will ask for it separately and say here how to take it back.

## 6.4 Your right to object

### Right to object under Art. 21 GDPR

**You have the right to object at any time, on grounds relating to your particular situation, to processing of your personal data that we carry out on the basis of Art. 6(1)(f) GDPR** — that is, the processing in this notice whose legal basis is our legitimate interest. That is: the server logs in § 2.1, the memory and personalisation in § 2.8, the activity notifications in § 2.10, the error monitoring and AI tracing in § 2.11, the device storage and the processing that goes with it in § 2.12, the correspondence in § 2.13, the

findability of public forum content in § 2.14, the evidence records and the optional cancellation question in § 2.16, the automatic linking of a Slack identity in § 2.18, and the sign-up and sign-in checks in § 5. Where we rely on § 2.16 to establish or defend a legal claim, an objection may not be enough on its own — the exception in the next paragraph is written for exactly that case, and we will say so rather than simply not act.

If you object, we will no longer process the data concerned **unless we can demonstrate compelling legitimate grounds that override your interests, rights and freedoms, or the processing serves to establish, exercise or defend legal claims.**

**Where we process your data for direct marketing, you may object at any time and we will stop. No reasons are needed and no balancing takes place.**

*For example:* if you object to the personalisation in § 2.8, we delete the memory held for your account and for your chats and stop deriving new memory for you — today a manual step we take on request rather than a setting you can flip; the assistant then answers each request without knowing anything about your earlier ones. Nothing else about your account changes. To object, write to [privacy@uprelic.com](mailto:privacy@uprelic.com) and say what you are objecting to.

## 6.5 Deleting your account

You can delete your account yourself, in **Settings** → **Account**. Deletion is not instant: the request starts a **21-day period during which you can cancel it**, after which it is carried out. We then delete your chats and messages, your files and the storage they occupy, your projects, your artifacts, your automations, your memory profiles, your integration connections and your search-index entries.

Some things survive, and you should know which:

- **Billing and accounting records**, because the law does not let us delete them for eight years (§ 2.9).
- **The account record itself**, which is kept but **stripped of the data that identifies you** — name, handle, picture, memory — with sign-in permanently disabled. That is what makes it possible to delete you without destroying other people's data: a project someone shared with you, a comment thread you took part in, an invoice we must keep.
- **The agreement and declaration records** described in § 2.16 — which version of the terms you accepted and when, and any cancellation or withdrawal you sent us, including the IP address and browser recorded with it. These exist to prove what was agreed, so deleting them with the account would defeat their purpose.
- **Public forum content** — posts and comments you made in project forums stay where they are, as described in § 2.14, though they are no longer attributed to a functioning account.
- **The link to your Slack identity**, if you ever made one — the workspace and user id, and the name and email address taken from your Slack profile (§ 2.18). This bullet is not a position

we are defending. It is a gap in the deletion job, and it is written here because you are entitled to know what happens rather than what was meant to happen.

**Content you shared with other people stays with them** — a chat, a project or an artifact you shared does not disappear from the workspaces of the people you shared it with. If you want that removed too, say so during deletion or write to [privacy@uprelic.com](mailto:privacy@uprelic.com).

Disconnecting is worth doing separately: deleting your account removes our record of an integration, but you should also revoke Uprelic's access at the provider itself.

## 6.6 Getting a copy of your data

You can request a full export yourself, in **Settings** → **Account**. It is prepared in the background; when it is ready we email you download links. The bundle carries your profile and the declarations we hold for you, your chats and their messages (with the names of the files you attached), the projects you own with a list of the files and links you added to them, your forum posts and comments, and your account memory and your chat memories, as JSON files in one or more **ZIP archives** — the bundle is split when it is large. It does not carry the files themselves, your artifacts or your skills, because you can download each of those in the product — a file from its menu, an artifact from the artifact view, a skill as a ZIP of its whole folder. Two things are in neither the bundle nor behind a download button: **your automations**, and **the memory held with a project together with your work profile** (§ 2.8). Write to [privacy@uprelic.com](mailto:privacy@uprelic.com) and we will prepare either.

`/signed-url/route.ts:70-77` returns a null signed URL and `ArtifactShell.tsx:129-134` hands the browser a 0-byte file; the Library kebab fails loudly on the same case. `Apdf_doc` exposes only the rendered PDF and its page images, never the JSON document model that is the actual source (`web/app/api/share/artifact/[artifactId]/pages/route.ts:77-100``). Both are filed, as GitHub #30 and GitHub #31. Fix them and the sentence stands as written; until then it overstates the artifact download by two types.]

**The download links are valid for 48 hours.** After they expire, the prepared copy is deleted from our servers — which is deliberate: an export gathers a great deal about you into one place and should not sit on a link indefinitely. You can request a new one at any time.

If you would rather have the export another way, or need data the self-service export does not cover, write to [privacy@uprelic.com](mailto:privacy@uprelic.com) and we will prepare it manually.

## 6.7 Complaining to a supervisory authority

You may complain to a data protection supervisory authority, in particular in the Member State of your habitual residence, your place of work, or the place of the alleged infringement (Art. 77 GDPR). The authority responsible for us is:

**Berliner Beauftragte für Datenschutz und Informationsfreiheit** Alt-Moabit 59–61, 10555 Berlin, Germany Telephone: +49 30 13889-0 Email: [mailbox@datenschutz-berlin.de](mailto:mailbox@datenschutz-berlin.de)

You are welcome to raise the matter with us first — it is usually faster — but you do not have to.

## 6.8 Other self-service tools

- **Cancel a contract:** [uprelic.com/cancel-contracts](https://uprelic.com/cancel-contracts) (§ 312k BGB).
- **Withdraw from a contract:** [uprelic.com/withdraw](https://uprelic.com/withdraw) (§ 356a BGB).

Both store the record described in § 2.16.

---

## § 7 Children

---

**Uprelic is not for children under 16.** You must be at least 16 years old to create an account (Art. 8 GDPR as applied in Germany, § 3.1 of the Terms). We do not knowingly process the personal data of children under 16. If you believe a child has created an account, write to [privacy@uprelic.com](mailto:privacy@uprelic.com) and we will delete it.

---

## § 8 Changes to this notice

---

We update this notice when the product changes or the law does. **This version takes effect on 30 July 2026.**

Where a change is material — a new category of recipient, a new purpose, a transfer to a country we did not previously use — **we will tell you before it takes effect**, by email or by a notice in the product, and not only by quietly replacing this page.

**Previous versions remain available.** Every version of this notice that has been in force can be retrieved from the version selector on this page, in both languages, with the date it took effect. You should not have to take our word for what this notice used to say.