Data Processing Agreement

Effective as of 14 November 2026

Sessionly Data Processing Agreement. The current version is published at https://sessionly.dev/dpa.

IN SHORT

- When your organization uploads its engineers' sessions to the hosted Sessionly service, your organization is the controller and I, Kristiyan Ivanov, am its processor. This agreement sets out the terms Article 28 GDPR requires for that.
- It forms part of the Terms of Service and applies automatically when you use the Hosted Service. No signature is needed; if your records need a countersigned copy, email me and I will return one.
- The Hosted Service is not open yet. It will run the published code, so what it stores is already known and described here; where it runs will be added to Annex III before the first org is created.
- Full prompts, file contents, command output and source code are never uploaded. What is stored is a record of each session, by default without any text, with the security measures in Annex II.
- Sub-processors are listed in Annex III. I will tell you 30 days before adding one, and you may object.
- I will tell you about a personal data breach affecting your data without undue delay and in any case within 48 hours of becoming aware of it.
- Local use and a self-hosted server are not covered: I process nothing for you there.

This Data Processing Agreement (the "DPA") is between the organization that holds an account on the hosted Sessionly service ("you" or the "Customer") and Kristiyan Ivanov, a self-employed individual registered in Bulgaria (BULSTAT 180490275), of ul. Vasil Petleshkov 31-33, apt. 17, 1510 Sofia, Bulgaria ("I," "me" or the "Processor"). It supplements the Terms of Service (the "Terms") and governs the processing of personal data that you, your machines, workflows and team members upload to the Hosted Service. Capitalized terms not defined here have the meaning given in the Terms.

How this DPA is concluded. This DPA forms part of the Terms. It applies automatically, without signature, from the moment an org is created for you on the Hosted Service or a

session is first uploaded to it, whichever is earlier, and it binds the organization the org was created for. If your organization needs a signed copy for its records, email kristiyan@sessionly.dev with your organization's legal name and address, and I will return a countersigned copy of the current version. A signed copy does not change its terms.

1. Definitions
2. Scope and roles
3. Details of the processing
4. My obligations as processor
5. Sub-processors
6. International transfers
7. Assistance
8. Personal data breaches
9. Deletion and return
10. Audits and information
11. Your obligations
12. Liability
13. Term, precedence and governing law
Annex I: Details of processing
Annex II: Technical and organisational measures
Annex III: Sub-processors

---

1. Definitions

- **"Data Protection Law"** means Regulation (EU) 2016/679 (the "GDPR"), the GDPR as it forms part of the law of the United Kingdom (the "UK GDPR") and the UK Data Protection Act 2018, the Swiss Federal Act on Data Protection, and any national law implementing or supplementing them, in each case as they apply to the processing under this DPA.
- **"Customer Data"** means personal data that you, your machines, workflows or team members upload to the Hosted Service, as described in Annex I. It includes the session records your engineers' machines upload, the pull request reports the Sessionly GitHub Action sends, and the names, email addresses and roles you attach to access tokens.
- **"Hosted Service"** has the meaning given in the Terms: the hosted Sessionly service, including its dashboard and API, and the `sessionly` command-line program and GitHub Action when they talk to it.
- **"Sub-processor"** means a third party I engage to process Customer Data on my behalf.
- **"Personal data," "controller," "processor," "data subject," "processing," "personal data breach"** and **"supervisory authority"** have the meaning given in the GDPR.

2. Scope and roles

2.1 You are the controller, I am the processor

For Customer Data, you are the controller (or a processor acting for your own customers, in which case you warrant that your instructions to me are covered by their instructions to you) and I am your processor. I process Customer Data only to provide the Hosted Service to you under the Terms and this DPA.

2.2 What this DPA does not cover

- **Data for which I am the controller.** The email address and name you use to open an account or order a License, billing data and support correspondence are processed by me as controller, as described in the Privacy Policy. Where the same piece of data serves both purposes (for example a team member's email address, which I need to recognise their token and which you entered to give them access to your org), this DPA governs my processing of it on your instructions and the Privacy Policy governs my own.
- **What stays on your engineers' machines.** The `sessionly` program keeps its raw capture, its ledger and its configuration on the machine it runs on. That data, which includes full prompts and file contents, never reaches me, and I do not process it.
- **Coding agents and GitHub.** The makers of the coding agents your engineers use, and GitHub, process your data under your agreements with them. They are not my Sub-processors, and I do not control what they do. The GitHub Action runs in your GitHub account, on your instruction.
- **Self-hosted installs.** If you run the Sessionly server on your own infrastructure, I do not process any data for you and this DPA does not apply. A Business or Enterprise License is verified offline and sends nothing to me.

3. Details of the processing

The subject matter, duration, nature and purpose of the processing, the categories of data subjects and the categories of personal data are set out in Annex I. In outline: the `sessionly` program on each engineer's machine builds a record of each coding agent session, applies the content level chosen on that machine, and uploads the record; the GitHub Action reports pull requests as hashes of the lines they add; the Hosted Service

stores both, matches them, and shows totals and individual sessions to the people in your org whose tokens allow it. Full prompts, file contents, command output and source code are never uploaded.

---

4. My obligations as processor

4.1 Instructions

I will process Customer Data only on your documented instructions, including with regard to transfers to a third country, unless EU or Member State law to which I am subject requires otherwise; in that case I will tell you about the legal requirement before processing, unless that law prohibits it on important grounds of public interest. Your documented instructions are: the Terms; this DPA; the settings you and your engineers make in the dashboard, the API, the command-line program and the GitHub Action (tokens and roles, the content level, line hashes, which repositories report pull requests); the sessions and reports your machines and workflows upload; and any further written instruction you email me and I agree to. If I believe an instruction infringes Data Protection Law, I will tell you without delay and may pause the processing it concerns until it is resolved.

4.2 Confidentiality

I run the Hosted Service myself. No employee or contractor of mine has access to Customer Data. If I ever give anyone access, they will first be bound by a written confidentiality obligation at least as protective as this DPA, and where they act as a Sub-processor they will be added to Annex III under Section 5.

4.3 Security

I will implement and maintain the technical and organisational measures in Annex II, and at least those required by Article 32 GDPR, taking into account the state of the art, the costs of implementation, and the nature, scope, context and purposes of the processing and the risks to data subjects. I may update Annex II as the software changes, provided the overall level of protection does not fall. The current measures are also described on the security page, and the code that implements them is public.

4.4 Records and cooperation

I keep the record of processing activities that Article 30(2) GDPR requires of a processor and will cooperate, on request, with the supervisory authority competent for you in the performance of its tasks.

4.5 Only what is needed

I will not use Customer Data for any purpose other than providing the Hosted Service to you: not to train models, not to build profiles, not to assess the people in it, not for advertising, and not for marketing to third parties. I may inspect metadata to operate the Hosted Service, enforce the Terms and keep the Hosted Service secure, as Section 4.3 of the Terms says.

---

5. Sub-processors

5.1 General authorisation

You give me general written authorisation to engage the Sub-processors listed in Annex III, and to replace or add Sub-processors under Section 5.2. Annex III states what each Sub-processor does and what it sees; it matches the service providers listed in the Privacy Policy.

5.2 Changes and objections

Before I add or replace a Sub-processor that will process Customer Data, I will email the admins of your org at least 30 days in advance and update Annex III. You may object within those 30 days on reasonable, documented data protection grounds. If you do, we will discuss the concern in good faith; if I cannot address it within a further 30 days (for example by not using the Sub-processor for your org), either of us may terminate the affected org's use of the Hosted Service on written notice, and I will refund any prepaid fees for the period after termination. Continued use of the Hosted Service after the 30-day notice period without an objection is acceptance of the change.

5.3 Flow-down and responsibility

I engage each Sub-processor under a written contract that imposes data protection obligations at least as protective, in substance, as those in this DPA, in particular regarding security. Where a Sub-processor fails to fulfil its data protection obligations, I remain fully liable to you for its performance, as Article 28(4) GDPR provides.

6. International transfers

I am based in Bulgaria. The Hosted Service is not open yet; the country and region it runs in will be stated in Annex III before the first org is created.

I will not transfer Customer Data outside the European Economic Area, the United Kingdom or Switzerland, or allow a Sub-processor to, except as described in Annex III and in compliance with Chapter V GDPR (and the equivalent provisions of the UK GDPR and Swiss law where they apply). For each Sub-processor that processes Customer Data in a country without an adequacy decision, I rely on an appropriate safeguard offered by that Sub-processor: the EU Standard Contractual Clauses adopted by the European Commission (Decision (EU) 2021/914), Module 3 (processor to processor), with the UK International Data Transfer Addendum and the Swiss adaptations where they apply, or, where the Sub-processor is certified under the EU-US Data Privacy Framework and its UK and Swiss extensions, that certification. I check which applies when I engage a Sub-processor and when its terms change; copies of or links to the relevant clauses or certification are available on request. If a safeguard I rely on is invalidated, I will work with you in good faith to put an alternative in place or stop the transfer.

7. Assistance

7.1 Data subject requests

Taking into account the nature of the processing, I will assist you with appropriate technical and organisational measures, insofar as this is possible, in responding to requests from data subjects exercising their rights under Chapter III GDPR (access, rectification, erasure, restriction, portability, objection). Some of this you can do yourself: an engineer can see exactly what their machine uploads with `sessionly export`, an admin can read any session and pull request report through the dashboard and the API, as JSON, and can revoke tokens. The Hosted Service does not yet let you delete a single session yourself; email me and I will delete it. For that and anything else the tools do not cover, I will help within a reasonable time, normally within 10 business days. If a data subject contacts me directly about Customer Data, I will not respond beyond acknowledging the request and pointing them to you, unless the law requires me to, and I will forward the request to you without undue delay.

7.2 Security, breaches, impact assessments and prior consultation

Taking into account the nature of the processing and the information available to me, I will assist you in meeting your obligations under Articles 32 to 36 GDPR: security of processing, notification of personal data breaches to the supervisory authority and to data subjects, data protection impact assessments, and prior consultation. Collecting records of employees' work may call for an impact assessment on your side. The information in this DPA, its Annexes, the security page and the public source code will usually be what you need for one; if you need more, email me.

7.3 Costs

Assistance under this Section is free of charge unless your requests are manifestly excessive or repetitive, in which case I may charge a reasonable fee that we agree in advance.

---

8. Personal data breaches

If I become aware of a personal data breach affecting Customer Data, I will notify the admins of your org by email without undue delay, and in any case within 48 hours of becoming aware of it. The notification will describe, to the extent then known, the nature of the breach, the categories and approximate number of data subjects and records concerned, the likely consequences, the measures taken or proposed to address it and mitigate its effects, and a point of contact. Where I cannot provide all of this at once, I will provide it in phases without undue further delay. I will cooperate with you and take the reasonable steps you ask of me to investigate, contain and remedy the breach. My notification is not an admission of fault or liability.

A personal data breach means a breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, Customer Data. Unsuccessful attempts (such as requests with an unknown token or port scans) are not breaches and will not be notified.

9. Deletion and return

9.1 During the term

- **Sessions and pull request reports** are kept while your org exists, unless your Order states a retention period. Uploading a session again replaces the earlier copy, so lowering the content level on a machine and uploading again removes the previews of the sessions that machine still has.
- **Tokens** stop working when an admin revokes them. A revoked token's entry (email address, name, role and dates, never the token) is kept so that your org can see who had access.
- **Deletion on request.** I will delete particular sessions, a person's sessions or a repository's reports when an admin of your org asks me to by email.
- **Export.** At any time an admin can read every session and report through the API as JSON.

9.2 At the end

When the Terms end for your org, or earlier if you ask me in writing to delete the org, I will delete all Customer Data within a reasonable time, normally within 30 days, including sessions, pull request reports and tokens, and I will confirm the deletion by email on request. Two exceptions: data I must keep under EU or Member State law is kept only as long as that law requires and remains protected by this DPA; and copies in database backups persist until the backups rotate, during which they are used only to recover the service. Export what you need before closing your account, because after deletion I cannot return it. The data on your engineers' machines is not affected.

10. Audits and information

I will make available to you all information necessary to demonstrate compliance with Article 28 GDPR and this DPA, and allow for and contribute to audits, including inspections, conducted by you or an auditor you mandate. Because the Hosted Service is run by one person on a hosting provider's infrastructure, the following is how that works in practice:

- **Documentation first.** This DPA, its Annexes, the security page, the Privacy Policy and the public source code at github.com/Klvanow/sessionly describe what the Hosted Service does and how it is protected. Sessionly has no SOC 2 or ISO 27001 report, and I do not claim one.
- **Questionnaires on request.** I will answer a reasonable written security or data protection questionnaire once in any 12-month period, and more often after a personal data breach affecting your Customer Data or when a supervisory authority requires it. I will answer within 30 days of receiving it.
- **On-site inspections** only where Data Protection Law or a supervisory authority requires one that documentation and a questionnaire cannot satisfy. They take place on at least 30 days' written notice, during business hours, no more than once a year, at your cost, under a confidentiality agreement, and must not disturb other customers' data or the service. I cannot give access to a hosting provider's data centres; for those you rely on the provider's own compliance documentation, which I will help you obtain.

If an audit finds that I am not complying with this DPA, I will fix the problem without undue delay and tell you what I changed.

---

11. Your obligations

- You are responsible for the lawfulness of the Customer Data you upload to the Hosted Service, including having a legal basis, and for your instructions to me being lawful.
- Session records describe the work of your engineers. You are responsible for informing them, before their sessions are collected, of what is collected, what is uploaded, who can see it and what it is used for; for any information or consultation of employees or their representatives that the law requires where you operate; and for carrying out a data protection impact assessment where one is required.
- You decide the content level. With previews on, the start of every prompt and command is uploaded and may contain personal data about anyone, or a secret; keep them off unless you need them. The Hosted Service is not designed for special categories of personal data (Article 9 GDPR) or data about criminal convictions; do not upload such data.
- You decide who gets an admin token, and therefore who sees every engineer's sessions, and you revoke tokens when people leave.

- You keep the admins' email addresses on your org current, since that is where I send Sub-processor notices and breach notifications, and you tell me without delay if you believe a notice has not reached you.
- You remain responsible for your agreements with the makers of your coding agents and with GitHub.

---

12. Liability

Each party's liability arising out of or relating to this DPA, whether in contract, tort or otherwise, is subject to the exclusions and limitations in Section 13 (Limitation of Liability) of the Terms, and the amounts under this DPA and the Terms together count towards the single cap in Section 13.2 of the Terms. Nothing in this DPA limits either party's liability where Data Protection Law does not allow it to be limited, or affects a data subject's rights under Article 82 GDPR.

---

13. Term, precedence and governing law

13.1 Term

This DPA takes effect when the Terms do for your org and remains in force for as long as I process Customer Data for you, including the deletion period in Section 9.2.

13.2 Precedence

For matters of data protection, this DPA takes precedence over the Terms and any other agreement between us to the extent they conflict. Where a Sub-processor's Standard Contractual Clauses apply to a transfer and conflict with this DPA, the Clauses prevail for that transfer. Everything else about our relationship, including fees, warranties, liability and termination, is governed by the Terms.

13.3 Changes

I may update this DPA where Data Protection Law, a supervisory authority, a change in Sub-processors or a change in the Hosted Service requires it, in the way the "Changes" paragraph of the Terms allows: by posting the new version on the website with a new effective date and, for material changes, emailing the admins of your org at least 30 days before it takes effect. Changes that reduce protection for data subjects require

your agreement, which may be given by continuing to use the Hosted Service after the notice period.

13.4 Governing law and venue

This DPA and any dispute arising out of or relating to it are governed by the laws of the Republic of Bulgaria and subject to the exclusive jurisdiction of the courts of Sofia, Bulgaria, as Section 16 of the Terms provides, without prejudice to any right a data subject or supervisory authority has under Data Protection Law. Where Data Protection Law requires that a provision of this DPA be governed by the law of an EU Member State, Bulgarian law applies.

13.5 Severability and notices

If any provision of this DPA is held invalid, the rest remains in force and the invalid provision is replaced by one that comes closest to its purpose and complies with Article 28 GDPR. Notices under this DPA go to the admins of your org by email, and to me at kristiyan@sessionly.dev or by post to the address at the top of this DPA.

Annex I: Details of processing

Parties

- **Controller (data exporter):** the Customer, the organization holding the org on the Hosted Service, acting through its admins.
- **Processor (data importer):** Kristiyan Ivanov, self-employed individual registered in Bulgaria (BULSTAT 180490275), ul. Vasil Petleshkov 31-33, apt. 17, 1510 Sofia, Bulgaria, kristiyan@sessionly.dev.

Subject matter

Operation of the Hosted Service for the Customer: a server that receives a record of each AI coding agent session from the Customer's engineers' machines and reports of the Customer's pull requests, and shows what the sessions cost, what they changed and how much of it survived, with the dashboard and API around it.

Duration

For as long as the Customer holds an org on the Hosted Service, plus the deletion period in Section 9.2 of the DPA.

Nature of the processing

- Receiving the session records the Customer's machines upload and the pull request reports its workflows send; checking each against the published schema; storing it, filed under the owner of the token that sent it.
- Matching the hashed lines of a pull request against the hashed lines of sessions by the people who committed to it; summing sessions per person, repository, model, agent and day; showing sessions, totals and pull requests in the dashboard and the API to the people whose tokens allow it.
- Hosting, backing up and deleting the above.

Purpose

To show the Customer what its coding agent sessions were, what they cost and how much of what agents wrote survived, per person, repository, model and agent, and otherwise to provide the Hosted Service as described in the Terms. No other purpose.

Categories of data subjects

- **The Customer's engineers:** employees, contractors and others whose machines upload sessions to the Customer's org.
- **People with access:** the admins, members and others the Customer issues tokens to.
- **Pull request authors:** people who open pull requests in repositories where the Customer runs the GitHub Action, whether or not they use Sessionly.
- **Anyone mentioned in a preview**, if the Customer uploads previews: a prompt or command can name or describe any person.

Categories of personal data

- **Never uploaded:** full prompts, the agent's replies, file contents, the output of commands, source code, and absolute paths under a home directory. They stay on the engineer's machine.

- **Session records (always):** a random machine identifier; the agent and session identifier; the git user name and email address and the repository's remote address, as configured on the machine; start and end times, active time and the timing of each turn; counts of prompts, tool calls, subagents and permission requests; for each tool call its tool, status, duration and output size; the files read or changed, as paths relative to the repository or home directory, with lines added and removed; models, token counts and cost; the rewrite measurement per file. Together these show when and on what a person worked with an agent.
- **Previews (only if switched on):** the first 200 characters of each prompt; a one-line summary of each tool call (shell command, search pattern or URL, cut to 200 characters); the first 200 characters of an error; the session's title; SHA-256 hashes of prompts and tool inputs. Free text, which may contain any personal data the engineer typed.
- **Line hashes (only if switched on):** per changed file and model, the first 8 bytes of the SHA-256 of each line the agent added.
- **Pull request reports:** repository name, pull request number, title and address; the author's GitHub user name; state, branch names and dates; counts of commits, reviews, requests for changes, review comments by people and by bots, and review rounds; per changed file, hashes of the added lines. The email addresses on the pull request's commits are received with each report, used to select the sessions to compare with, and not stored.
- **Tokens:** the email address (stored lower-case), optional name and role each token belongs to; the token's SHA-256 hash; when it was created, last used and revoked.

Special categories of data

None intended. The Hosted Service is not designed for special categories of personal data or data about criminal convictions; the Customer must not upload such data (Section 11 of the DPA).

Frequency, retention and location

- **Frequency:** continuous: a connected machine uploads new and changed sessions when a turn ends (at most every five minutes) and when a session ends, and the GitHub Action reports when a pull request is opened, pushed to, reviewed or closed.
- **Retention:** sessions, pull request reports and tokens while the org exists, unless the Order states a shorter period or the Customer asks for deletion earlier.

- **Location:** to be stated in Annex III before the first org is created on the Hosted Service.

Annex II: Technical and organisational measures

These are the measures in place as of the effective date, written from the code, which is public at github.com/Klvanow/sessionly. The security page describes them in more detail. Measures that depend on where the Hosted Service runs (transport encryption in front of it, network isolation, backups) will be added here before it opens. Sessionly has no SOC 2 or ISO 27001 certification.

Data minimisation, before anything is uploaded

- Full prompts, file contents and command output stay in the raw capture on the engineer's machine, in files and directories readable only by that user (modes 0600 and 0700). The session record holds metadata and 200-character previews, not content.
- Before upload, file paths are made relative to the repository or to the home directory, and the content level is applied. The default level, `metadata`, removes all text from prompts, commands, search patterns, URLs, errors and titles, and all content hashes. Previews and line hashes are each a separate, deliberate switch on the machine.
- `sessionly export` prints the exact payload for any session, so an engineer can see what would be sent before connecting to a server. The payload's JSON Schema is published.
- Pull requests and sessions are compared as hashes of lines. No source code is sent to the Hosted Service by the program or by the GitHub Action.

Encryption in transit

- The `sessionly` program refuses to connect to a server over plain HTTP unless the server is on the same machine, so uploads to the Hosted Service always use HTTPS.

Access control and authentication

- Access is by bearer token. Tokens are generated from 24 random bytes, shown once when created, and stored only as SHA-256 hashes. They can be revoked at any time; when each was last used is recorded. There are no passwords.
- Each token has one role. A member reads only the sessions filed under their own email address. An admin reads the whole org and manages tokens. A ci token can report pull requests and check that it is accepted, and is refused everything else.

- An upload is filed under the owner of the token that sent it, whatever name the record carries, and a session cannot be overwritten by a different person.
- On the engineer's machine the token is kept in a configuration file readable only by that user and is never printed.
- Operationally, one person, the Processor, administers the Hosted Service. No employees or contractors have access to production systems or Customer Data.

Integrity

- Every upload is validated against the published schema and every pull request report is checked field by field before it is stored; one that does not match is refused. Uploads are limited to 32 MB.
- The capture hook on the engineer's machine only observes: it cannot block or change an agent's session, and events that would change the agent's behaviour when hooked are not registered.

Platform

- The server does not write a log of requests, and the program and the server send no telemetry to me or to anyone else. License files are verified offline.
- The website is static, loads no scripts, and is served with a Content Security Policy of `default-src 'none'` with only its own stylesheet, fonts and images allowed, plus `X-Content-Type-Options`, `X-Frame-Options` and `Referrer-Policy` headers.

Incident handling and transparency

- Personal data breaches affecting Customer Data are notified to the Customer without undue delay and in any case within 48 hours of my becoming aware of them (Section 8 of the DPA).
- Security problems can be reported to kristiyan@sessionly.dev, which I read myself. The whole codebase, including the capture program, the server and the GitHub Action, is public, so the measures above can be verified by reading it.

Annex III: Sub-processors

The Hosted Service is not open yet, so no Sub-processor processes Customer Data as of the effective date. Before the first org is created, this Annex will name the provider that hosts the Hosted Service and its database, with its location and transfer basis, and any

other provider that will handle Customer Data. The makers of the Customer's coding agents and GitHub are not Sub-processors (Section 2.2 of the DPA).

Sub-processor	Purpose	What they see	Location and transfer basis
Vercel	Hosting this website	Web request logs for sessionly.dev. No Customer Data	United States company; serves the website from its edge network worldwide. Transfer relies on Vercel's Standard Contractual Clauses or Data Privacy Framework certification.

Vercel is listed so that the list here and the one in the Privacy Policy are the same. Changes to this list are notified under Section 5.2 of the DPA.