

Data Processing Agreement

Rflect AG

Version 3.0 - 20 October 2026

1. Subject matter of the agreement

1.1 Description. Rflect AG, with its registered office in Zurich, Switzerland, postal address c/o Ivan Jovanovic, Neunbrunnenstrasse 118, 8050 Zürich, registered in the Commercial Register of the Canton of Zurich under company number CHE-402.686.433 (the "Processor"), provides a software-as-a-service solution to support educational processes to the customer (the "Controller").

1.2 Preamble. This agreement sets out the data-protection framework and the obligations for handling personal data in accordance with Regulation (EU) 2016/679 (GDPR), the Swiss Federal Act on Data Protection (FADP), the cantonal data-protection laws applicable to the Controller where relevant, and other applicable data-protection laws (together, "Data Protection Law").

1.3 Relationship. The Controller is the data controller for all personal data of Authorised Users (students, lecturers, teaching assistants, coaches and other natural persons granted access to the Services by the Controller) processed through the Services. The Processor acts exclusively as data processor on behalf of the Controller and processes personal data only on the Controller's documented instructions. This agreement does not apply to personal data that the Processor processes as an independent controller for its own purposes (account administration of Institution contacts, billing, customer-relationship management and sales communication, and the operation of its website); that processing is described in the Processor's Data Protection Notice.

1.4 Scope. The Processor provides services under the Institutional Terms of Service and the Order Form (together, the "Contracts"). The scope, purpose and duration of the processing are defined in the Contracts and, where applicable, in supplementary written instructions of the Controller.

1.5 Term. This agreement applies for the term of the Contracts and survives until the Processor has returned or deleted all personal data processed on the Controller's behalf in accordance with Section 11.

1.6 Order of precedence. In the event of conflict between this agreement and the Contracts on matters of personal-data processing, this agreement prevails.

2. Nature, purpose and categories of processing

2.1 Nature and purpose. The Processor processes personal data to operate the Rflect platform, to enable reflective learning activities, to provide insights, feedback suggestions and assessment support to lecturers based on shared reflections, and to deliver AI-supported educational features to Authorised Users where they have opted in. The AI-supported features and the data each of them processes are described in Annex 4.

2.2 Categories of data subjects. Authorised Users (students, lecturers, teaching assistants, coaches, administrators) of the Controller.

2.3 Categories of personal data.

- Account data: email address, name, surname and related identifiers.

- Self-assessment data: gender and age range (where collected for academic-study purposes).
- Institutional affiliation: program, class, cohort and related context.
- Reflection content submitted by Authorised Users. The Services are not designed or intended to collect special categories of personal data (Art. 9 GDPR / Art. 5 lit. c FADP), and the parties take the measures in Section 2.7 to keep such data out of the Services. The Processor recognises that free-text content may nonetheless contain data on the intimate sphere, on health or otherwise particularly sensitive personal data within the meaning of Art. 5 lit. c FADP / Art. 9 GDPR, and that AI-supported analysis of reflection content may amount to profiling (Art. 5 lit. f FADP / Art. 4(4) GDPR) and, depending on the program design, to high-risk profiling (Art. 5 lit. g FADP). Such data is treated as unsolicited and handled in accordance with Section 2.7.
- Assessment and feedback data: feedback and assessment suggestions generated for lecturers on shared reflections, and the lecturer's own assessments, where the Controller uses these features.
- Usage logs and technical metadata generated by use of the Services.

2.4 Sharing model. For program reflections, the lecturer sets the sharing level of each question, for example private, or shared with the lecturer or the class, by name or anonymously. The student sees the level before answering; it cannot be changed once students have started. The Processor implements this sharing level as both a technical and a legal access boundary, as further set out in Section 2.5. Authorised Users are informed that AI-supported features may be applied to shared reflections on the Controller's instruction (Section 2.6(a)).

2.5 Private reflections - standing instruction and carve-out. The Controller acknowledges that reflections submitted under a "private" sharing level are, by design and by contract, not accessible to the Controller, its lecturers, its coaches or its administrators. The following standing arrangements apply, irrevocably, for the term of this agreement and beyond:

(a) The Controller waives, in advance, any right to (i) access reflections marked "private", (ii) instruct the Processor to disclose them or to surface them in dashboards, lecturer or coach views, exports or reports, or (iii) require aggregate statistics that would allow re-identification of private content.

(b) This standing instruction may not be amended, suspended or overridden by side letter, change order, oral arrangement or any amendment of this agreement, except by a written amendment that simultaneously and transparently re-consents the affected Authorised Users with prior in-app notice.

(c) On termination of the Contracts, private reflections are deleted or, at the option of the Authorised User, returned to that user and never to the Controller.

(d) This Section 2.5 is concluded as a genuine third-party stipulation in favour of Authorised Users within the meaning of Art. 112(2) CO: Authorised Users may invoke the Controller's waiver directly.

(e) If the Processor receives a binding court order, regulatory demand or criminal-procedure request that requires it to surface private content, the Processor will notify the Controller (and, where permitted, the affected Authorised User) without undue delay and will, to the extent legally permissible, delete the affected private reflections rather than release them.

(f) Audits permitted under Section 8 may verify the technical separation between private and shared reflection content, but may not be used to access private content.

(g) The Processor does not review private reflection content through its own personnel and does not surface private reflection content on Processor-internal dashboards. This commitment is subject only to Section 2.5(e) (binding court order, regulatory demand or criminal-procedure request).

(h) AI-supported features do not take decisions producing legal or similarly significant effects for Authorised Users. Where a feature supports the assessment of an Authorised User's work, it provides suggestions to the lecturer; the lecturer reviews them and takes every assessment decision. The Processor supports the Controller in meeting its obligations under Regulation (EU) 2024/1689 (AI Act) by providing the instructions for use, logging and human-oversight information required for the AI-supported features.

The Processor warrants in the technical and organisational measures (Annex 3) that private reflections are architecturally inaccessible to Controller-side dashboards and administrative roles.

2.6 AI-supported features. The Processor distinguishes between AI processing of shared reflections and AI processing of private reflections, and groups AI-supported features into the two data-flow categories and the three risk categories defined in Annex 4:

(a) **Shared reflections.** AI-supported features that operate on reflections shared with a lecturer or with a class, whether by name or anonymously (for example, AI-supported insights and assistance for lecturers, such as the "Teaching Assistant" feature) are provided under this agreement on the Controller's general instruction. No separate Authorised User opt-in is required, for a reflection submitted under the sharing level shown before answering. The shared reflection content included in the prompt is limited to what the lecturer is permitted to access. Where attribution is permitted, answer-reading tools may include the student's name or label and account ID; answers provided anonymously exclude stable author identifiers. Authorised roster lookups and student-analysis tools may also include a student's name, email address and account ID (Annex 4, category B).

(b) **Private reflections.** AI processing of reflections marked "private" (for example, the "Go Deeper" feature, providing personalised feedback to the Authorised User on their own reflection) takes place only when the Authorised User actively requests it for that reflection. The first time the Authorised User uses such a feature, the Services inform the user clearly and unambiguously that the reflection text is processed by an AI model and that each request constitutes consent to that processing. The Controller has no role in obtaining this consent; it is collected and managed by the Processor on a user-by-user basis on the Controller's standing instruction set out in this Section. Each request constitutes the Authorised User's express consent (Art. 6(7) lit. a and b FADP / Art. 9(2)(a) and Art. 7 GDPR) on which the Controller relies; the Processor records the date, time and version of the notice shown for each request and makes these records available to the Controller on request, so that the Controller can demonstrate consent (Art. 7(1) GDPR). The Services remain fully usable without such requests. The Controller acknowledges that it cannot consent on behalf of an Authorised User to AI processing of private reflections, in particular where the content is particularly sensitive personal data. No AI processing of private reflections takes place without a request. The prompt for such features contains the text the user has written and no identifier (Annex 4, category A).

(c) **No model training.** Reflection content as well as lecturer created content is not used to train AI models, whether by the Processor or by sub-processors. This does not apply to anonymised or synthetic material that has been verified as anonymous in accordance with Annex 5 and no longer constitutes personal data under Section 3.3(c).

(d) **AI inference processing.** AI inference for AI-supported features runs through EU inference endpoints operated by the AI inference sub-processor identified in Annex 2. To reduce latency and cost, the request prompt may be held in the sub-processor's ephemeral, encrypted prompt cache for up to one (1) hour from last use, after which it is automatically deleted. The cache stays within the EU inference environment, is not accessible to the Processor's personnel, and is never used to train AI models (Section 2.6(c)). The Teaching Assistant's code sandbox (AWS Bedrock AgentCore Code

Interpreter, Frankfurt) holds working copies of the data it analyses for at most 30 minutes per session. No input or output is otherwise retained by the sub-processor beyond the inference request.

(e) **No analysis of health or wellbeing.** The Processor does not detect, score, flag or monitor the health, mental state or wellbeing of Authorised Users, does not provide crisis detection or alerting, and instructs its AI-supported features to decline requests for such analysis. The Services are a teaching tool and not a medical, therapeutic or crisis service.

(f) **Feature governance.** New or changed AI-supported features are classified into the risk categories in Annex 4. Category 1 changes may be introduced at any time and are listed in release notes. Category 2 changes are announced to the Controller's administrators in release notes when they are enabled; the Controller may object within thirty (30) days on reasonable data-protection grounds, in which case Section 5.2, third sentence, applies by analogy. Category 3 changes require thirty (30) days' prior written notice to the Controller, a joint review with the Controller if the Controller so requests, and an update of Annex 4 before they are enabled; the Controller may object on reasonable data-protection grounds, in which case Section 5.2, third sentence, applies by analogy.

(g) **Model changes.** The Processor may change the AI models used where this serves the purpose of the Services. Model changes are listed in release notes. A change of AI provider or of the inference region is a change of sub-processor under Section 5.2.

2.7 Special categories of personal data.

(a) **Design.** The Services are not designed to elicit special categories of personal data. The Processor does not provide reflection-question templates that invite such content and provides the Controller with guidance on question design.

(b) **Notice to users.** When a student starts using Rflect, the Services display a notice that Rflect is a learning tool and not a medical, therapeutic or crisis service, and that health details about the student or anyone else should not be entered.

(c) **Controller instructions.** The Controller instructs its Authorised Users not to enter special categories of personal data into the Services and instructs its lecturers, when designing reflection questions, not to elicit such data and to delete or report such content when they encounter it. Where the Controller deliberately designs a program to elicit such content, the Controller is responsible for a legal ground under Art. 9(2) GDPR / Art. 31 FADP and informs the Processor in advance.

(d) **Report and removal.** A lecturer can report a shared reflection that contains special categories of personal data. The Processor removes the reported content from the production system and deletes every Teaching Assistant conversation in which the reported reflection was included (individual messages within a conversation are not edited), in each case within five (5) working days of the report, records the removal with metadata only (date, program, reason), and informs the affected Authorised User that the reflection can be resubmitted without such content. Removed content leaves the backups when the applicable backup retention in Annex 3 expires. The Controller may also instruct the Processor to remove such content under Section 7.

(e) **Unsolicited data.** Special categories of personal data that reach the Services despite the measures in this Section are unsolicited. The Processor does not process them for any purpose other than storage, display to the persons the Authorised User has shared the reflection with, the AI-supported features under Section 2.6 on the same terms as any other reflection, and removal under (d), and the technical and organisational measures in Annex 3 apply to them in full.

3. Processor's obligations

3.1 Documented instructions. The Processor processes personal data only on the Controller's documented instructions, including with regard to transfers to a third country. The Contracts and this agreement constitute the initial instructions. The Controller may amend or supplement instructions in writing or text form at any time. Where the Processor is required by Swiss or EU law to process data otherwise, it will inform the Controller of that legal requirement before processing, unless the law prohibits such notice.

3.2 Confidentiality. The Processor ensures that persons authorised to process personal data have committed themselves to confidentiality or are under an appropriate statutory obligation of confidentiality, and have been instructed in the relevant data-protection requirements before processing.

3.3 Use of personal data by the Processor.

(a) The Processor processes personal data only to provide, operate, secure and support the Services on the Controller's behalf and on the Controller's documented instructions.

(b) For the operation, security and support of the Services, the Processor may review shared reflection content (reflections submitted under a sharing level, set by the lecturer, that shares them with the lecturer or with the class, whether by name or anonymously) and outputs generated by Rflect features on the basis of such shared content (for example, conversations with the Teaching Assistant) where this is necessary to investigate a user-reported or Controller-reported issue, to maintain the security and safe operation of the Services, and to ensure and improve their proper functioning.

(c) The Processor may use anonymised or aggregated data derived from shared reflections, provided such data no longer constitutes personal data, for performance reporting, academic research and publications, and for improving and further developing the Services. This includes training, fine-tuning and evaluating AI models that are used exclusively within the Services. Such material is not used to train, fine-tune or otherwise adapt models of third-party AI providers, such as Anthropic or OpenAI. To evaluate models, the Processor may send such material as ordinary inference requests to the AI models available through the sub-processors in Annex 2 (Section 2.6(d)); neither the material nor the results are used to train those models.

This Section constitutes the Controller's documented instruction within the meaning of Art. 28(3)(a) GDPR. On this instruction, and for the Controller's purpose of ensuring that no personal data of its data subjects is used for the purposes set out in this Section, the Processor may generate synthetic material from shared reflections and their educational context (reflection questions, learning outcomes and lecturer feedback) as a method of anonymisation. Private reflections, special categories of personal data within the meaning of Art. 9 GDPR, particularly sensitive personal data within the meaning of Art. 5 lit. c FADP, and content reported under Section 2.7(d) are excluded. To exclude such data, the Processor may check shared reflections automatically; the check serves exclusively this exclusion, no result other than an exclusion flag is stored, and the flag is not used for any other purpose. Reflections that an Authorised User has deleted or has ceased to share are excluded from all subsequent generation runs; this does not affect material already anonymised. Synthetic material is generated from anonymised reflections. Only material verified as anonymous under the method documented in Annex 5 may be used for the purposes in this Section. The Processor may update Annex 5 where the level of protection is maintained or improved, notifies the Controller of material updates and makes the current version available on request. The Processor reviews Annex 5 against applicable regulatory guidance, including final EDPB guidance on anonymisation, and adapts it where required.

Such material, and models trained on it, are not made available to third parties as datasets or published. The Controller may exclude its data from this processing at any time by written notice, with

effect for future processing. The Processor implements the exclusion without undue delay, at the latest within 30 days of receipt; material already verified as anonymous remains unaffected. Unless it has excluded its data, the Controller informs its Authorised Users of this processing in its own data protection information; the Processor provides a text module for this purpose.

(d) Customer support is part of the Services. Support requests and conversations with the Processor's support assistant that contain personal data of Authorised Users are processed on the Controller's instruction under this agreement. The Processor retains them for 365 days from the last exchange, so that issues can be traced and escalations handled, and deletes them thereafter and in any event in accordance with Section 11 on termination of the Contracts. Contact and contract data of the Controller's staff, ticket metadata and billing data are processed by the Processor as independent controller under Section 1.3. Support content may be used to improve the support or to automate it only as anonymised or synthetic material, under the conditions of Section 3.3(c) and the method in Annex 5.

3.4 Data-subject rights assistance. The Processor assists the Controller, taking into account the nature of the processing and the information available to the Processor, in fulfilling its obligations under Arts. 12–22 GDPR / Arts. 25-29 and 32 FADP, including responding to requests for access, rectification, erasure, restriction, portability and objection. If an Authorised User exercises rights by contacting the Processor directly, the Processor informs the Controller without undue delay and awaits the Controller's instructions before taking action, except where the Processor is required by Data Protection Law to act on the request itself.

3.5 Compliance assistance. The Processor assists the Controller in complying with its obligations under Arts. 32–36 GDPR / Arts. 7-8 and 22-24 FADP, including security measures, breach notification, data protection impact assessments and prior consultation, taking into account the nature of the processing and the information available to the Processor.

3.6 Data portability and account migration.

(a) **Standing authorisation.** The Controller authorises the Processor, as a standing instruction granted once under this agreement, to let any Authorised User export or migrate the personal data that the user has themselves authored, on that user's own request and without further approval from the Controller.

(b) **Scope.** This covers reflections and other content authored by the Authorised User - irrespective of sharing setting (private, shared with lecturer, or shared with class) - together with feedback directed specifically to that user. It does not extend to content authored by or about other data subjects, including other students' reflections, class or group threads, and cohort-level data.

(c) **Copy, not removal.** Export or migration is provided as a copy or transmission to a destination of the user's choice. It does not delete or alter the Controller's own records, which the Controller retains.

(d) **Migration window.** Where the Controller removes or de-provisions an Authorised User's account, the Processor holds deletion of that user's authored data for thirty (30) days, during which the user may exercise the right in (a). After that window, the Processor completes deletion in accordance with Section 11.

4. Technical and organisational measures

4.1 The Processor implements appropriate technical and organisational measures to ensure a level of security appropriate to the risk, in accordance with Art. 32 GDPR / Art. 8 FADP. The measures in force at the date of this agreement are summarised in Annex 3 (Technical and Organisational Measures) and described in full in the Processor's TOM document, which the Processor makes available to the Controller.

4.2 Backup and retention. The Processor maintains a tiered backup regime (daily incremental, daily, weekly and monthly snapshots). Backup retention periods are specified in Annex 3. The Processor operates a documented restore process under which deletion requests received between the snapshot date and any restore date are re-applied by an authorised operator as a required step of the documented restore procedure before service resumes, so that data deleted at the request of the Controller or an Authorised User is not re-introduced into the production system by a restore.

4.3 Material changes. The Processor reviews the technical and organisational measures regularly and notifies the Controller of material changes that affect the security of the processing. The Processor updates this agreement, its annexes and the TOM document when a release changes the categories of personal data processed, the AI processing described in Annex 4, the sub-processors or the retention periods.

5. Sub-processors

5.1 General authorisation. The Controller authorises the engagement of the sub-processors listed in Annex 2. The Processor maintains the up-to-date list of sub-processors and the description of their processing.

5.2 Change notice and objection. The Processor notifies the Controller in writing at least **30 days** before engaging a new sub-processor or replacing an existing one. The Controller may object to the change on reasonable data-protection grounds within that 30-day period. If the Processor cannot accommodate the objection, the Controller may terminate the Contracts in respect of the affected Services on reasonable notice, without penalty.

5.3 Equivalent obligations. The Processor binds each sub-processor by written contract to data-protection obligations substantially equivalent to those set out in this agreement, in particular with regard to confidentiality, security measures, sub-processing, international transfers and assistance with data-subject rights.

5.4 International transfers. Where a sub-processor processes personal data outside Switzerland and the EEA, the Processor implements an appropriate transfer mechanism (EU Standard Contractual Clauses (Commission Implementing Decision (EU) 2021/914) with the adaptations for Swiss law recognised by the FDPIC, reliance on the EU-US and/or Swiss-US Data Privacy Framework where the recipient is certified for the relevant track, or another mechanism recognised under Data Protection Law) and conducts a transfer-impact assessment that it makes available to the Controller on request.

5.5 Review. The Processor reviews its sub-processors at least once a year against documented criteria (data-protection terms, security certifications, processing location and transfer mechanism, incident history, continued necessity), records the result with a date, and makes the record available to the Controller on request.

6. Breach notification

6.1 Notification. The Processor notifies the Controller of any confirmed or suspected personal-data breach without undue delay and in any event within **24 hours** of becoming aware of it, providing at least:

- a description of the nature of the breach, including, where possible, the categories and approximate number of data subjects and records concerned;
- the name and contact details of the Processor's data-protection contact;
- a description of the likely consequences of the breach; and

- a description of the measures taken or proposed to address the breach and to mitigate its possible adverse effects.

6.2 Mitigation. The Processor takes all reasonable measures without undue delay to contain the breach and to mitigate its consequences for data subjects and for the Controller.

6.3 Third-party events. If the Controller's data at the Processor are jeopardised by attachment, seizure, insolvency proceedings or any other third-party measures, the Processor informs the Controller without undue delay, unless prohibited by court order or law. The Processor informs the relevant authorities and counterparties that decision-making authority over the data lies exclusively with the Controller.

7. Right to issue instructions

7.1 The Processor processes personal data only within the scope of the Contracts and on the Controller's instructions, subject to Section 3.1.

7.2 Authorised persons. Instructions may be given by the persons identified in Annex 1 (Authorised Persons). For the Controller, all active lecturer accounts and the contact identified by the Controller as its data-protection contact are authorised to issue instructions in relation to data they administer.

7.3 Unlawful instructions. If the Processor considers that an instruction infringes Data Protection Law, it notifies the Controller without undue delay. The Processor may suspend the execution of the instruction until the Controller has confirmed or modified it, and may decline to execute a manifestly unlawful instruction.

8. Audit rights

8.1 The Controller is entitled to verify the Processor's compliance with this agreement and with Data Protection Law within an appropriate scope. The Controller exercises its audit rights only to the extent required to meet its legal obligations and in a manner that does not unduly disrupt the Processor's operations.

8.2 The Processor provides the Controller, within a reasonable deadline, with the information and documentation required to evidence compliance, including independent audit reports (such as ISO 27001) where available. On-site audits may be carried out on reasonable notice.

8.3 Audit rights are subject to the limitations in Section 2.5(f) regarding private reflections.

9. Liability

Liability is governed by the applicable Data Protection Law and, where the Contracts contain a more specific provision, by the Contracts. Liability for intent, gross negligence, personal injury and fraud is not excluded (Art. 100(1) CO; Art. 101(3) CO). Nothing in the Contracts limits either party's liability towards data subjects under applicable Data Protection Law (in particular Art. 82 GDPR and Art. 28 ff. CC) or the parties' statutory rights of recourse for compensation paid to data subjects (Art. 82(5) GDPR).

10. Extraordinary right of termination

The Controller may terminate the Contracts without notice in full or in part if the Processor materially breaches this agreement, intentionally or with gross negligence violates Data Protection Law, or is unable or unwilling to execute a lawful instruction of the Controller. For minor (non-intentional, non-grossly negligent) breaches, the Controller grants the Processor a reasonable cure period.

11. Termination

11.1 Return or deletion. On termination of the Contracts and at the option of the Controller, the Processor returns or deletes all personal data processed on the Controller's behalf, and deletes existing copies, unless retention is required by Swiss or EU law. The Processor maintains documentation of the deletion and provides a certificate of deletion on request.

11.2 Exception for private reflections. Section 2.5(c) applies to private reflections: these are deleted or returned to the Authorised User at the user's option and then deleted, never to the Controller.

11.3 Surviving confidentiality. The confidentiality obligations in Section 3.2 survive termination.

12. Concluding terms

12.1 No retention beyond term. The Processor has no right to retain personal data under this agreement beyond the Contracts' duration unless required by law.

12.2 Amendments. The Processor may amend this agreement, including its Annexes, by written notice to the Controller's data-protection contact (Annex 1) or, if none has been named, to the Controller's contact addresses on file. Amendments take effect on the date stated in the notice, which may be the date of the notice. The notice summarises the changes. Amendments may not reduce the level of protection of personal data under this agreement. The Controller may object on reasonable data-protection grounds within thirty (30) days of the notice; in that case the Processor either continues to apply the previous provision to the Controller or, if this is not reasonably possible, the Controller may terminate the Contracts without penalty and receives a pro-rata refund of prepaid fees. The engagement of new sub-processors follows Section 5.2 and changes in risk category 3 follow Section 2.6(f); these require prior notice. Sections 2.5, 2.6(b) and 2.6(c) may be amended only as provided in Section 2.5(b). Amendments agreed individually between the parties require text form, including email or electronic signature.

12.3 Relationship to Contracts. The Contracts remain in force unless inconsistent with this agreement, in which case this agreement prevails on matters of personal-data processing.

12.4 Severability. If any part of this agreement is or becomes invalid, the remaining parts remain in force. The parties replace the invalid term with one that best reflects the original economic purpose.

12.5 Governing law. Swiss substantive law, together with the material Union law applicable to the processing, including the GDPR.

12.6 Jurisdiction. Exclusive jurisdiction lies with the competent court of Zurich, Switzerland.

Annex 1 — Authorised persons

On the part of the Processor (by function, not by name):

- Data-Protection Contact, reachable at privacy@rflect.ch.
- Customer Support, reachable at support@rflect.ch.

On the part of the Controller:

- All active lecturer accounts of the Controller.
 - The data-protection contact identified by the Controller in the Order Form.
-

Annex 2 — Approved sub-processors

Name	Description of processing	Location of processing	Third-country transfer
AWS	Infrastructure-as-a-service: hosting of the application, the database, logs, backups and all in-transit data through AWS cloud services, including email delivery (SES), push notifications (SNS) and speech-to-text for voice memos (Transcribe; the transcript is written to the Processor's own storage and removed once read, and the finished job is deleted from Transcribe), CloudFront and WAF, and AWS Chatbot for operational alerts. Additionally, the Processor uses AWS Bedrock AI inference to deliver AI-supported features, with Bedrock Guardrails and the AgentCore Code Interpreter in Frankfurt (working copies in sandbox sessions of at most 30 minutes). Where students opt in, private reflection content (text only, without name or email address) may be processed through AWS Bedrock to generate personalised feedback and insights. Students cannot opt out of processing of shared reflections by AWS Bedrock. All AI inference processing occurs within the EU. To reduce latency and cost, AWS Bedrock may hold the request prompt in an ephemeral prompt cache within the EU inference environment for up to one (1) hour from last use, after which it is automatically deleted; AWS Bedrock does not otherwise store data beyond the inference request and does not use customer data for model training.	EU	None by design (EU data region); any remote access by AWS personnel from outside the EU/Switzerland is governed by the AWS DPA (SCCs / DPF)
Sentry	Application logs and error diagnostics; no session replay or session recordings, request bodies are not captured, personal data is scrubbed	EU / Frankfurt	None
Pipedrive (CRM)	CRM and onboarding workflows. Receives the lecturer's name and email address, program name and dates, and sign-up and invitation events, transmitted via API. No student personal data and no reflection content.	EU (Pipedrive's EU data centres)	None
Google Workspace	Business email and customer support, including transmission of exports requested by the Controller. Processes personal data only where it is contained in an email or attachment; not a system of record.	EU (Workspace data region)	EU Standard Contractual Clauses under the Google Cloud DPA; the EU data-region setting covers stored data, not every processing step

Name	Description of processing	Location of processing	Third-country transfer
Apple (APNs) and Google (FCM), via AWS SNS	Delivery of push notifications to mobile devices. The notification text contains topic and program names and, for the daily action reminder, the text of the Authorised User's own action. No reflection content, no name and no email address.	USA (Apple and Google notification services)	EU Standard Contractual Clauses and Swiss-US Data Privacy Framework certification of the provider, under the Apple Developer Program License Agreement and the Google Firebase Data Processing and Security Terms
Slack (Salesforce)	Internal operations channel of the Processor. Receives automated email delivery alerts: complaint alerts containing only a count and a message identifier. No reflection content.	USA (Slack Pro workspace; no EU data residency option on this plan)	EU-US and Swiss-US Data Privacy Framework (Salesforce) and EU Standard Contractual Clauses under the Slack Data Processing Addendum
Thinking Machines Lab, Inc. (Tinker)	Managed infrastructure for training, fine-tuning and evaluating AI models used exclusively within the Services. Receives only anonymised or synthetic material verified as anonymous in accordance with Annex 5, together with related model checkpoints and technical training data. No personal data of Authorised Users, non-anonymised reflection content or re-identification mappings are transferred. Customer content is not used to train the provider's own models.	USA; subcontractors may process globally.	No transfer of personal data for this activity; only material verified as anonymous under Annex 5.

Sub-processor data-protection terms:

- AWS — Data Processing Addendum: <https://d1.awsstatic.com/legal/aws-dpa/aws-dpa.pdf>
- Sentry — Data Processing Addendum: <https://sentry.io/legal/dpa/>
- Pipedrive — Data Processing Addendum: <https://www.pipedrive.com/en/privacy/dpa>
- Google Workspace — Cloud Data Processing Addendum: <https://cloud.google.com/terms/data-processing-addendum>
- Slack — Data Processing Addendum: <https://slack.com/terms-of-service/data-processing>
- Google (FCM) — Firebase Data Processing and Security Terms: <https://firebase.google.com/terms/data-processing-terms>
- Apple (APNs) — Apple Developer Program License Agreement: <https://developer.apple.com/support/terms/>
- Thinking Machines Lab (Tinker) — Service Terms of Use, including Exhibit B (Data Processing Agreement): <https://thinkingmachines.ai/legal/terms/>

Annex 3 — Technical and organisational measures (summary)

- **Architectural separation of private reflections.** Private reflection content is stored and access-controlled such that it is not retrievable through Controller-side or Processor-internal dashboards, administrative roles or reporting endpoints, and is not surfaced in AI workflows other than the user's opted-in features.
- **Access control.** Access to production systems and internal administration tools is restricted to authorised Processor personnel bound by confidentiality (Section 3.2); administrative access to production systems is restricted to designated company accounts protected by multi-factor authentication; application requests pass a central authorisation check.
- **Encryption.** Encryption in transit (TLS) and at rest (AES-256, with keys held in AWS KMS or managed by AWS) for production data; encryption of backups.
- **Logging and monitoring.** Centralised logging, application error monitoring, security alerting. Logs contain metadata only and never reflection content, prompts or AI outputs; retention 90 days, applied uniformly for all Controllers.
- **Backups.** Daily incremental and full database backups, with weekly and monthly snapshots. Retention: point-in-time recovery 35 days, daily 90 days, weekly 365 days, monthly 730 days, with encrypted copies in a second EU region. A documented restore process re-applies deletions received between the snapshot date and the restore date. That the database backups in both EU regions can be restored is verified at least quarterly through automated restore tests, each of which is recorded; the re-application of deletions, the restore of files and point-in-time recovery are tested by hand on a quarterly checklist.
- **Vulnerability management.** Automated dependency and container-image scanning (weekly and on every build); findings remediated under documented targets by severity.
- **AI guardrails.** System-prompt boundaries and classifier-based filtering against special categories of data and against health or wellbeing analysis (Annex 4).
- **Sub-processor security.** Contractual confidentiality, data-protection and security commitments from each sub-processor as set out in Section 5; annual documented review (Section 5.5).
- **Incident response.** Documented incident-response procedures, including breach notification within the timeframes in Section 6.

Annex 4 — AI-supported features, data flows and risk categories

A. Feature categories by data flow

Category	Features	What enters the prompt	Output and storage	Opt-in and model
A. Text only	For example, Go Deeper (follow-up questions and feedback on the user's own reflection). Triggered by the student.	The text of the reflection or goal, the reflection question and, where set, the program's Go Deeper instruction. No name, email address or other identifier.	Shown to the student and stored with the reflection; follows its retention.	Opt-in required for private reflections (Section 2.6(b)). LLM on AWS, EU.

Category	Features	What enters the prompt	Output and storage	Opt-in and model
B. With identifiers	For example, the Teaching Assistant (assistance for the lecturer on shared reflections, including drafting feedback and assessment suggestions against the Controller's learning outcomes or rubric, for the lecturer's review). Triggered by the lecturer.	Shared reflections the lecturer may access, the lecturer's question and, for assessment support, the Controller's learning outcomes or rubric. Attributed answer reads may include a student's name or label and account ID; anonymous answer reads exclude stable author identifiers. Authorised roster and student-analysis lookups may additionally include name, email address and account ID.	Shown to the lecturer; conversation history stored in the Processor's database in the EU, deleted upon deletion of the author's account or on the Controller's request.	No separate opt-in (Section 2.6(a)); students are informed when they start using Rflect and when they share a reflection. LLM on AWS, EU.

Inference for both categories runs on AWS Bedrock EU inference endpoints (Section 2.6(d)); no content is used for model training (Section 2.6(c)); no wellbeing, health or crisis analysis is performed (Section 2.6(e)).

B. Risk categories for new or changed AI-supported features

Category	Definition	Notice to the Controller
1	Any change to a feature within category A or category B that adds no new category of personal data to a prompt, no new recipient and no new sub-processor.	None in advance; listed in release notes.
2	A new AI-supported capability offered to lecturers on shared content within the purpose in Section 2.1 (for example, support for the lecturer's own assessment under Section 2.6(h)) that adds no new category of personal data to a prompt, no new recipient and no new sub-processor.	Release notes to the Controller's administrators at enablement; the Controller may object within 30 days on reasonable data-protection grounds (Section 2.6(f)).
3	A new category of personal data in a prompt, AI processing of private reflections without the user's opt-in, a decision with legal or similarly significant effect taken by the AI, or a new AI provider or inference region.	30 days' prior written notice, joint review at the Controller's request, update of this Annex before enablement; objection right under Section 2.6(f). Where required, a fundamental rights impact assessment is supported before release.

C. Guardrails

The system prompts of all AI-supported features instruct the models not to generate questions aimed at special categories of personal data, to decline requests for health or wellbeing analysis of individual users, and to refer the requesting user to the usage policy. A classifier-based filter on inputs and outputs enforces the same boundary independently of the prompt.

D. Updates to this Annex

This Annex is updated when a feature moves between the categories in Part A, when a new category of personal data enters a prompt, when the AI provider or inference region changes, or when the retention of AI outputs changes. Part A names features by example; adding a feature to an existing category does not change this Annex. A change of AI model within the same provider does not change this Annex and is communicated under Section 2.6(g). Releases that do not touch these points do not change this Annex.

Annex 5 — Anonymisation and synthetic data method

This Annex sets out the anonymisation method under Section 3.3(c). It applies to shared reflections and, under Section 3.3(d), to support content.

A. Source data. The Processor uses reflections shared at extraction, their teaching context and support content, on the Controller's instruction. Private, deleted or reported records, records of Controllers that have excluded their data and records excluded by the Contracts are not used. Later changes to sharing settings apply to future anonymisation runs. Existing anonymised data is not deleted retroactively.

B. Sensitive data. Records containing data covered by Arts. 9–10 GDPR or Art. 5 lit. c FADP are identified automatically and excluded entirely; they are not cleaned up or rewritten. Apart from the exclusion flag, no result of this check is stored, and the flag is used for no other purpose.

C. Anonymisation. A large language model (LLM) removes, generalises or replaces identifying facts with generic or fictional details while preserving teaching value. Rare or unusual details that could single out a person are generalised or removed. In addition to names and places, the field of study, setting and course of events are changed where this is not necessary for teaching value. Fictional tasks may include consistent fictional data, for example fictional programmes, students and reflections, so that multi-step tasks can be tested. Anonymity is assessed against course participants and lecturers, Processor personnel and users of the trained models.

D. Verification. Each example passes the following checks:

(a) AI content check: AI checks every example against its sources for identifying facts, including content that has been reworded but, alone or in combination with other details, could still single out a person, including to someone who knows the course, such as a lecturer or fellow student. Passages and categories found are verified against the source. Examples with confirmed findings are rewritten and checked again, or rejected.

(b) Word matching: each example is also compared against all sources, ignoring case and treating punctuation as word boundaries. Examples with 12 or more consecutive matching words are rejected or rewritten and checked again.

(c) Human review: a person trained for this task reviews a random sample of each release. The size of the sample depends on the volume of the release and on the results of the automated checks.

Before use, the Processor validates the anonymisation method and the anonymity of the resulting data and models, including tests of whether a trained model reproduces source content or reveals whether a person's data was used for training. The thresholds for these tests and for the checks in (a) to (c) are documented by the Processor and made available to the Controller on request. Changes may only maintain or improve the level of protection; material changes are notified to the Controller. Only examples that have passed all checks are used. Each release is documented with metadata only.

Where material previously verified as anonymous is subsequently confirmed to contain personal data, the Processor excludes the affected examples and any identified derivatives containing that personal data from all future training, fine-tuning and evaluation runs.

E. Use and retention. Processing of personal data under this Annex, including anonymisation and verification, takes place in the EU through the sub-processors in Annex 2. Only outputs verified as anonymous, which no longer constitute personal data, may be processed outside the EU, including by providers not listed in Annex 2, for training, fine-tuning and evaluation of AI models used exclusively within the Services. Such providers act solely on the Processor's instructions. Datasets and model weights are not published or used by third parties for their own purposes, and the restrictions in Section 3.3(c) continue to apply. Access to source data and training data is restricted to personnel who need it for this processing.

Individual erasure requests apply to remaining personal data. They do not require tracing or deleting anonymous datasets or retraining anonymous models.

Temporary sources, intermediates, mappings and unreleased outputs are deleted after the checks, within seven (7) days of extraction. Intermediaries, mappings and unreleased outputs are not retained in backups beyond the standard rotation period of 90 days.

F. Review. The method is reviewed at least annually, after any confirmed finding in released material and when final EDPB guidance on anonymisation is published.