

Rflect AG

Technical and Organisational Measures (TOM)

pursuant to Art. 32 GDPR and Art. 8 revFADP in conjunction with Art. 1-6 DSV

Version 1.1, 20 October 2026

Provider	Rflect AG, Neunbrunnenstrasse 118, 8050 Zurich, Switzerland
Contact	Niels Rot, Co-Founder, info@rflect.ch
Hosting	AWS, Frankfurt region (eu-central-1), with encrypted backup copies in Ireland (eu-west-1). AI inference exclusively in EU regions via AWS Bedrock EU inference profiles.
Security owner	Benjamin Nater (Engineering Lead) is the security owner and coordinates vulnerability management and the sub-processor review (see sections 4 and 9).

This document describes the technical and organisational measures that Rflect AG, acting as data processor (Art. 28 GDPR / Art. 9 revFADP), has implemented to ensure data security. The infrastructure used (AWS) is certified to ISO/IEC 27001, 27017, 27018, 27701, 42001, SOC 2 Type 2 and BSI C5 Type 2; AWS Bedrock falls within the scope of these certifications.

1. Confidentiality

Area	Measure
Physical access	Pure cloud architecture with no own data centres. Physical security of the AWS data centres in Frankfurt (24/7 monitoring, multi-level access control, biometrics) per AWS SOC 2.
Authentication	End users sign in with email and password (Argon2id, with a salt per password; passwords stored under the earlier scrypt scheme are re-hashed with Argon2id at the next sign-in) or with a one-time code sent by email; through their learning management system via LTI 1.3 (OIDC login initiation and a signed ID token); or, where the institution has enabled it, through single sign-on with the institution's OpenID Connect identity provider or SWITCH edu-ID. Single sign-on only signs in existing accounts; with SWITCH edu-ID, only addresses in the institution's registered email domains are accepted. Sign-in and one-time code requests are rate-limited per email address and per IP address; password reset requests are rate-limited per IP address. Administrators: mandatory multi-factor authentication for AWS, the source repository and deployment pipeline, error monitoring and Google Workspace administration. Administrative access to production systems and to the internal administration tool is restricted to designated company accounts protected by multi-factor authentication.

Area	Measure
Authorisation (RBAC)	Access is granted through organisation roles (administrator, member) and per-program roles (lecturer, student) on a need-to-know basis; lecturers see only the students in the groups assigned to them. Every application request passes a central authorisation check. Each question carries a sharing level that the lecturer sets when creating the topic and that the student sees before answering (private, shared with the lecturer, shared with the lecturer anonymously, shared with the class, or shared with the class anonymously); it cannot be changed once students have started. Lecturers never see the content of private answers; for those they see only whether and when the question was answered, its length and the time spent. Direct access to the production infrastructure is limited to two named administrators through AWS IAM Identity Center with short-lived credentials, and all AWS management activity is recorded by an organisation-wide CloudTrail trail. The application authenticates to the database with short-lived IAM credentials; the master credential is held in AWS Secrets Manager for emergency use only.
Hardware isolation	The application runs as containers on Amazon ECS on EC2 instances built on the AWS Nitro System. Exec access into containers is disabled, and hosts expose no SSH port and hold no stored keys. Operator access goes only through AWS Systems Manager Session Manager with short-lived credentials and is limited to the two administrators; the start of every session is recorded in CloudTrail. The AWS Nitro System was independently assessed by NCC Group in 2023 (public report).
Pseudonymisation	Product monitoring and research use aggregated metrics by user type and program; logs and telemetry use pseudonymous identifiers instead of names. Personal-level data is accessed only by authorised Rflect staff for support, troubleshooting and quality review, never for routine analysis of individuals. Web analytics is cookieless.

2. Integrity

Area	Measure
Encryption in transit	All transmissions over public networks (web application, API, LTI, mobile app) use TLS 1.2 or higher with modern cipher suites only, and HSTS is enforced. Database connections are TLS-encrypted with certificate verification. Outbound email is sent over TLS whenever the recipient's mail server supports it.
Encryption at rest	AES-256 for the database (RDS), file storage (S3) and backups. Integration credentials are held in access-restricted database tables; only the authorised integration runtime can read them.

Area	Measure
Input and audit control	Every API request is logged with timestamp, method, route, status, duration, a request ID and, for signed-in users, the account ID; IP addresses, user agents and request or response bodies are not logged, and tokens and email addresses are removed from URLs. Application logs contain metadata only: in every environment the logger accepts only identifiers, one-way fingerprints, counts, sizes, durations, fixed status values and error types, and has no field for free text or error messages, so reflection content, prompts, AI outputs, names and email addresses are not written to logs. Application logs are kept, encrypted, in Amazon CloudWatch Logs and a self-hosted log store (Grafana Loki) in the same EU region. Access logs of the load balancer and CDN, which contain IP addresses, are kept for security purposes. The web application firewall keeps no per-request logs for the application; it records aggregate counts per rule only. Error monitoring (Sentry, EU instance) receives only pseudonymous account IDs, route patterns and sanitised error types; request bodies, error message text and session recordings are not sent. Security-relevant operations are additionally recorded in the database: account deletion requests and approvals (timestamp, approving account), program deletions (timestamp, deleting account) and permission grants (granting account). Every deletion is also recorded within minutes in a write-once deletion ledger; its records are kept indefinitely and cannot be altered or removed for three years. Logs are kept for 90 days and traces for 30 days, uniformly for all controllers; access is limited to authorised personnel. On request, Rflect provides a controller with log extracts for the retention period, filtered to the account identifiers of its users and the identifiers of its programs.
Application security	Server-side input validation and size limits on every request; database access only through parameterised queries; output escaping by the frontend framework and server templates; session cookies are HttpOnly, Secure and SameSite=Lax (SameSite=None and partitioned only for sessions started inside a learning management system via LTI), state-changing requests additionally require an allowed origin and a CSRF token, and cross-origin requests are limited to Rflect domains. Database in a private VPC subnet without internet routing, not publicly accessible.
Threat detection	Amazon GuardDuty monitors database login activity and storage access patterns and alerts the operations team within minutes.
Operational monitoring	Logs, traces, frontend errors and performance data are collected on a self-hosted Grafana stack in the same AWS EU region. It records pseudonymous account and session identifiers and no names, email addresses, IP addresses or content.

3. Availability and resilience

Area	Measure
High availability	The database runs as RDS Multi-AZ with a synchronously mirrored standby and automatic failover. Application containers run on hosts spread across several Availability Zones; if a zone fails, the application restarts automatically in another zone. Capacity is adjusted through Infrastructure as Code.
Protective measures	AWS Shield Standard (Layer 3/4 DDoS), AWS Web Application Firewall with rate limiting, and technical alerts for defined error conditions. Failed or aborted backup and cross-region copy jobs raise automatic alerts.

Area	Measure
Backup and recovery	Daily automated database backups with point-in-time recovery (35 days), plus daily, weekly and monthly backups of the database and the file store (S3: voice memos, attachments, profile photos) in an AWS Backup vault protected by a vault lock (governance mode), with encrypted copies, including transaction logs, in a second EU region (Ireland, eu-west-1). File storage is versioned so that overwritten or deleted files can be restored for 35 days. Documented restore procedure in which deletions made between the backup date and the restore date are re-applied, so that data deleted at the request of the controller or users is not reintroduced into the production environment by a restore (cf. DPA §4.2). That the database backups in both regions can be restored is verified at least quarterly through automated restore tests; date, scope, duration and outcome of each test are recorded. The re-application of deletions, the restore of files and point-in-time recovery are tested by hand on a quarterly checklist. Infrastructure as Code enables reproducible rebuilds.
Recovery objectives	Rflect does not commit to binding RTO/RPO values, because they depend on the type of incident and on cloud-provider availability. Indicative, not contractual: zone failure, no data loss, restart within minutes; restore of an earlier state, data loss limited to minutes, restart within a few hours; loss of the entire primary region, database data loss limited to minutes and uploaded files up to 24 hours, restart within several hours to a day.

4. Procedures for regular review

Area	Measure
Data protection management	Responsibility with the management; operational implementation by Engineering and Operations. Processing for institutional customers is governed by the Data Processing Agreement, which forms part of the Institutional Terms of Service (Section 3).
Vulnerability and patch management	Dependencies are monitored continuously with GitHub Dependabot (security alerts and fix pull requests for Go, npm and Python; weekly update pull requests for GitHub Actions); the Go backend is scanned weekly with govulncheck, and every container image, including its base image, is scanned by Amazon ECR when it is pushed. Critical security updates can be rolled out at short notice. Findings are triaged by severity and tracked to closure with the following remediation targets: critical, mitigation within 24 hours and fix within 7 days; high, 14 days; medium, 30 days; low, next regular release.
Sub-processor review	All sub-processors are reviewed at least once a year by the security owner against documented criteria: data-protection terms in force, security certifications, processing location and transfer mechanism, incident history, and continued necessity. The result is dated and retained. Changes to the list follow DPA §5.
Code quality	Application code and database schema changes reach production only from a release pull request that must pass the full build, the automated test suites and a staging deployment; this gate has no bypass. Infrastructure changes are applied by authorised operators from a reviewed Terraform plan. Staging and production run in separate AWS accounts. Rflect policy prohibits placing production data or program exports in source repositories or giving them to AI development tools; automated checks block exported data files and real email addresses from being committed. AI development tools have no access to the production database or to user content; the operational telemetry they can read carries pseudonymous identifiers only.

5. Sub-processors

The database (PostgreSQL on AWS RDS) is located in eu-central-1 in a private VPC subnet. All primary data processing takes place within the EU. The authoritative list of sub-processors is Annex 2 of the Data Processing Agreement; this section mirrors it for the sub-processors that process personal data.

Provider	Location	Purpose	Transfer safeguard
AWS	EU, primarily Frankfurt; backup copies in Ireland	Compute (ECS), database (RDS), storage (S3), content delivery and firewall (CloudFront, WAF), AI inference for AI-supported features (Bedrock, EU inference profiles, Anthropic Claude models; not persisted beyond the one-hour prompt cache) with content filtering (Bedrock Guardrails), a code-execution sandbox for the Teaching Assistant (Bedrock AgentCore Code Interpreter, sessions of at most 30 minutes), email delivery (SES), push notifications (SNS), speech-to-text for voice memos (Transcribe; transcripts are not retained by AWS), internal alert delivery (AWS Chatbot)	Processing within the EU
Sentry	EU (de.sentry.io)	Error monitoring. Receives pseudonymous account IDs, route patterns and sanitised error types; URLs are scrubbed, and request bodies, error message text and session recordings are not sent	Processing within the EU
Pipedrive	EU (Pipedrive's EU data centres)	CRM. Rflect's backend calls the Pipedrive API directly. Pipedrive receives the name and email address of lecturers who register themselves, create a program or invite students, the program name and its start and end dates, and whether students were invited by email or by link. No student or invitee data and no reflection content.	Processing within the EU
Google Workspace	EU (Workspace data region)	Business email and customer support. Receives support requests submitted in the app (the requester's name, email address and message), account-deletion notices and email delivery alerts, and is used to send exports a controller has requested; exports made in the app are downloaded directly from the platform. Not a production system of record. For ordinary account administration, billing and sales communication Rflect acts as an independent controller.	Google Cloud Data Processing Addendum, including EU Standard Contractual Clauses for any processing outside the EU

Provider	Location	Purpose	Transfer safeguard
Apple (APNs) and Google (FCM), via AWS SNS	USA	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 user's own action; the payload also contains internal identifiers and a link into the app. No reflection content, no name and no email address.	EU Standard Contractual Clauses and Swiss-US Data Privacy Framework certification of the provider (Apple Developer Program License Agreement; Google Firebase Data Processing and Security Terms)
Slack (Salesforce)	USA (Slack Pro, no EU data residency)	Internal operations channels. They receive email delivery alerts: complaint alerts contain only the number of complaining recipients and a message identifier, without address or subject line. Infrastructure alerts contain no user content. No reflection content.	EU-US and Swiss-US Data Privacy Framework (Salesforce); EU Standard Contractual Clauses under the Slack DPA

6. Data breaches and incident response

Area	Measure
Incident register	Security incidents are recorded in a restricted internal incident register kept outside the source repository (time of awareness, severity, scope, affected controllers, assessment, decisions and measures). Affected controllers are informed in accordance with the process below.
Process	Containment, assessment, notification of the controllers without undue delay and in any case within 24 hours of becoming aware, root-cause analysis, remediation. Controllers are informed in good time so that they can meet their notification obligations under Art. 33 GDPR / Art. 24 revFADP.

7. AI-specific measures

Area	Measure
Models	Anthropic Claude models via AWS Bedrock, inference only. Rflect does not train, fine-tune or otherwise adapt these or any other third-party provider's models. No personal data is used to train any model; only material verified as anonymous under Annex 5 of the DPA may be used to train, fine-tune and evaluate own models used exclusively within Rflect (see Anonymisation and synthetic data below).

Area	Measure
Anonymisation and synthetic data	Used only on the controller's instruction under DPA §3.3(c); each institution can exclude its data at any time by written notice. Source: shared reflections, their teaching context and support content; private, deleted or reported records are not used. Records containing special categories of personal data are identified automatically and excluded entirely; apart from the exclusion flag, no result of this check is stored. A language model removes, generalises or replaces identifying facts with fictional details. Each example is compared against all sources (12 or more consecutive matching words trigger an AI content check), and the method and the anonymity of the resulting data and models are validated before use; only examples that pass all checks are used, and each release is logged without content. All processing of personal data in this process, including anonymisation and verification, takes place in the EU through the sub-processors in DPA Annex 2. Only material verified as anonymous may be processed outside the EU for training, fine-tuning and evaluation of models used exclusively within Rflect, by providers acting solely on Rflect's instructions; it is never used to train third-party providers' models, and datasets and model weights are not published. Temporary working data is deleted within seven days of extraction and kept out of logs and backups. The method is reviewed at least annually (DPA Annex 5).
Contractual guarantees	AWS Bedrock stores neither prompts nor completions. Anthropic has no operational access to Bedrock prompts or responses; the models run on AWS-managed infrastructure. Under the Bedrock Commercial Terms, Anthropic may not train models on customer data.
EU inference	All language-model inference runs via the EU inference profiles of AWS Bedrock, and voice memos are transcribed by Amazon Transcribe in Frankfurt; data remains within the AWS network inside the EU. To reduce latency and cost, the prompt of any AI-supported feature may be held in an ephemeral, encrypted prompt cache within the EU inference environment for up to one hour after its last use, after which it is automatically deleted. 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 prompts or responses are stored permanently by the provider.
AI features and data flows	AI features fall into two categories. Category A, text only: the prompt contains the text the user has written, the reflection question and, where the lecturer has set one, the program's Go Deeper instruction, and no name, email address or other identifier (Go Deeper). Category B, with identifiers (Teaching Assistant): the prompt contains shared reflections the lecturer may access. Attributed answer reads may include the student's name or label and account ID; anonymous answer reads exclude stable author identifiers. Authorised roster and student-analysis lookups may also include name, email address and account ID. Marking summaries and program insights process shared reflections without identifiers. Teaching Assistant conversation history is stored in Rflect's own database in the EU, not at the inference provider. The full description per feature is Annex 4 of the DPA.
Special categories of data	The platform is not designed to collect special categories of personal data. Students are told when they start using Rflect that it is a learning tool, not a medical, therapeutic or crisis service, and are asked not to enter health details about themselves or anyone else. Lecturers receive guidance on question design. A report function and a defined removal process exist for such content that is entered nonetheless (see section 8).

Area	Measure
Guardrails	System prompts instruct the models not to generate questions aimed at special categories of data, to decline requests for health or wellbeing analysis of individual users, and to refer to the usage policy. A classifier-based filter on inputs and outputs enforces the same boundary independently of the prompt.
Product boundary	Rflect does not detect, score, flag or monitor the health, mental state or wellbeing of users and has no crisis-detection or alerting function. It is a teaching tool; the student terms state that it is not a medical, therapeutic or crisis service.
Model and feature changes	A change of model version within the same provider is evaluated against a fixed evaluation set before release, logged, and listed in release notes. A change of AI provider or inference region is a sub-processor change (DPA §5.2, 30 days' notice). New AI features are classified into the risk categories in Annex 4 of the DPA: category 1 changes are listed in release notes; category 2 capabilities (for example assessment support for lecturers) are announced in release notes at enablement with a 30-day objection right; category 3 changes are notified 30 days in advance.
Student privacy	Rflect staff do not review private reflection content, and it is not shown on Rflect-internal dashboards; the only exception is a binding court order, regulatory demand or criminal-procedure request (DPA §2.5). AI processing of private reflections, including transcription of private voice memos, takes place only when the student requests it for that reflection. Research and product development use only anonymised or aggregated data (DPA §3.3(c)). AI-supported features for lecturers on shared content (e.g. Rflect TA) are part of the lecturer toolset; students are informed of this when they start using Rflect and when they share a reflection.

8. Retention and deletion

Area	Measure
Technical logs	90 days, encrypted, in the EU region (Amazon CloudWatch Logs and the self-hosted log store); traces 30 days. The same period applies to all controllers.
Reflections and shared content	For the term of the contract; deletion at the end of the contract or on request (DPA §11). Private reflections are never released to the institution.
AI outputs	Go Deeper feedback is stored with the reflection it belongs to and follows its retention. Teaching Assistant conversation history is stored in Rflect's database in the EU and is deleted when the lecturer deletes the conversation, when the lecturer's account is erased, or earlier on the controller's request. The prompt cache is deleted one hour after its last use.
Reported content	Content reported as containing special categories of data is removed from the live database within five working days; every Teaching Assistant conversation that contains the reported reflection is deleted in the same period (individual messages are not edited); a removal log with metadata only (date, program, reason) is kept for 24 months. Removed content disappears from backups when the backup retention expires.
Support requests	Support requests and conversations with the support assistant are retained for 365 days from the last exchange and deleted at the end of the contract at the latest (DPA §3.3(d)). Contact and ticket metadata of institution staff and billing data are kept by Rflect as independent controller.

Area	Measure
Contract / master data	Contract and accounting records are kept outside the platform for 10 years (statutory Swiss retention periods).
End of contract	At the end of the contract, productive institutional data is deleted. Backups are retired within their retention period (database point-in-time backups 35 days, daily backups 90 days, weekly backups 365 days, monthly backups 730 days; file versions 35 days). The same retention applies to the copies in the second EU region.
Right to erasure	Individual requests (Art. 17 GDPR / Art. 32 revFADP) are handled promptly. Private reflections are deleted at the end of the contract or, at the user's option, returned to the user, and never released to the institution.
Uploaded files	Deleting a file (voice memo, attachment, profile photo) removes it from the application immediately. The previous version is retained for 35 days for recovery from accidental deletion, then purged automatically; the file then remains only in backups until those expire (see End of contract). The restore procedure re-applies file deletions in the same way as database deletions, so a restore never reintroduces deleted files. Voice memos and their transcripts follow the same rule.
Data portability and account migration	Users can export or migrate the content they have authored (reflections regardless of sharing status, as well as feedback addressed to them) without the institution's consent. If the institution removes or deactivates an account, deletion of the self-authored data is held for 30 days so that users can export or migrate it (cf. DPA §3.6).

9. Organisational measures

Area	Measure
Confidentiality	Employees are subject to the statutory duty of loyalty, including confidentiality (Art. 321a CO).
Awareness	The data-protection and information-security rules that apply to the work (no personal data in logs, no production data outside the production environment, incident response) are written down and binding for every team member.
Departure	Immediate revocation of all access upon an employee's departure.
Supplier management	Sub-processors are selected according to data-protection and security criteria (EU hosting preferred, established providers with their own compliance programs). Providers' data-processing agreements (including EU Standard Contractual Clauses where applicable) apply. Annual documented review, see section 4.
Compliance documentation	Every compliance document, questionnaire answer and security statement is reviewed and released by a named person before it is sent to a controller. AI-assisted drafting and code scanning are used as tooling; statements are verified by that person against the systems described here.

These measures are reviewed regularly and adapted to technology and the threat landscape. The current version is available at any time from info@rfect.ch.