# Data Handling, Privacy & AI Overview

Aperon, Inc. · Public · October 2026

> This overview explains what data the Aperon Platform processes, where it is kept, how long it is kept, how it is returned and deleted, and how the platform uses AI.

## What data the platform processes

| Category | Examples |
|---|---|
| Trial communications | Trial-related email correspondence and attachments delivered to the platform |
| Trial documents | Documents that authorized users upload to a trial |
| Derived records | Trial facts, issues and timelines extracted from communications and documents, plus corrections made by users |
| Account data | Users' names, email addresses, roles, and authentication and audit events |

Depending on how a customer uses the platform, trial communications may include personal or health information. All customer data is classified as **Confidential** under our Data Management Policy. It is restricted to authorized personnel on an as-needed basis, encrypted at rest and in transit, and securely disposed of when no longer needed.

## Where data is kept

- Each customer's data is kept in that customer's **dedicated Google Cloud environment**: its own database, document storage and secrets. It is never mixed with another customer's data.
- Data is hosted in **Google Cloud regions in the United States** by default. Uploaded documents are stored redundantly across two US regions. Database backups stay in the same region as the database. Other hosting regions may be available by agreement.
- Customer records do not pass through Cloudflare. Cloudflare provides DNS, transactional account email and the status page.

## Data minimization in logs and email

- The bodies of ingested emails are never written to application logs. Logs carry only metadata.
- Emails sent by the platform, such as invitations, password resets and service notices, contain only links and account information, never clinical or trial data.
- Usage telemetry in the platform's software libraries is turned off.

## Retention

| Data | Retention |
|---|---|
| Customer data | For the term of the customer agreement |
| Application audit trail | For as long as the workspace exists. It cannot be modified or deleted. |
| Database backups | 30 days, with point-in-time recovery for the previous 7 days |
| Infrastructure audit logs | 7 years in a tamper-resistant archive |
| After termination | Deleted within 90 days unless the contract specifies otherwise |

## Export, return and deletion

- **Export at any time.** Authorized users can export a trial's complete record as JSON, CSV or a printable document, with a reconciliation manifest. They can also export the audit trail in a machine-readable format. Every export is recorded in the audit trail.
- **At the end of the contract**, by default, we give the customer a complete copy of their data and then permanently delete their dedicated environment: the database, document storage, secrets and cloud project. If the customer's contract requires records to be kept for a set period, an export is placed in tamper-resistant retention storage for that period instead.
- **Deletion requests** can be sent to privacy@aperontrials.com. We reply within two business days and complete deletion within 30 days of confirming the request, with written confirmation. Each request is tracked under a documented procedure that requires approval and keeps evidence of completion.
- **Backups** expire on their normal schedule and are never used to restore deleted data.

## Use of AI

The Aperon Platform uses **Google Gemini models through Google Cloud Vertex AI**, and Google Document AI, to:

- classify incoming trial communications;
- extract trial facts and flag anomalies for review;
- answer users' questions about a trial, citing the source records.

### How customer data is protected in AI processing

| Commitment | How it is met |
|---|---|
| Only Google Cloud processes data | AI calls are made from within the customer's dedicated Google Cloud environment, over Google's private network, under Google Cloud's enterprise terms. No other AI provider receives customer data. |
| No model training on customer data | Under Google Cloud's terms for Vertex AI, Google does not use customer data to train or fine-tune its models without permission. Aperon does not currently use customer data to train or fine-tune models. Any future use of de-identified data to improve models would be governed by customer agreements and disclosed in advance. |

| Commitment | How it is met |
|---|---|
| Scoped access | The AI sees only the trial the user has selected, through bounded, read-only, approved views of that trial's data. It cannot change records or act on the user's behalf. |
| Rules before models | Deterministic rules run before any model call, and only the content needed for a task is sent. |
| Prompt-injection defenses | Content from emails and documents is treated as untrusted. It is isolated within prompts and neutralized so that it cannot change the model's instructions. |
| Temporary processing data | Intermediate document-parsing output is deleted as soon as parsing finishes, with an automatic cleanup rule as a backstop. |

### Accuracy and human oversight

- Every factual sentence in an AI answer cites the source record it came from, and answers do not rely on general knowledge.
- The model version used for processing is recorded in the audit trail.
- People review AI output. Any change to a record is attributed to an authenticated user in the audit trail.

## Subprocessors

| Subprocessor | Purpose | Customer data |
|---|---|---|
| Google Cloud Platform | Hosting, database, document storage, secrets, and AI processing (Vertex AI, Document AI) | Yes, within each customer's dedicated environment |
| Cloudflare | DNS, transactional account email, status page | Account email addresses only; no trial records |
| GitHub | Source code management and CI/CD | No |

The current list is maintained on the Aperon Trust Center.

## Privacy contacts

- Privacy and data requests: privacy@aperontrials.com
- Security concerns: security@aperontrials.com