

Subprocessor List

1.1 · Customer-facing information

Subprocessor List

This page lists the subprocessors and processor-side provider paths Sefira currently uses, together with optional providers that are used only when the relevant feature, integration, or deployment configuration is enabled.

It should be read together with the [Privacy Notice](/privacy), the [Data Processing Addendum](/dpa), the [Security White Paper](/compliance/security-white-paper), and the [AI Act Readiness Note](/compliance/ai-act-readiness).

How to read this list

The most important distinction is:

1. the core product system of record is the application plus its PostgreSQL-compatible database layer; and
2. additional providers are used only when the relevant feature path requires them, such as storage, email, billing, AI, or customer-enabled integrations.

This means Customer Personal Data is not sent to every provider in every workflow. Exposure depends on enabled features, deployment choices, and connected integrations.

Current subprocessors and provider paths

Provider	Service area	Why it is used	Main data categories involved	Optionality	Region / transfer note
---	---	---	---	---	---
Neon	PostgreSQL database hosting	Core application system of record for workspace records, audit evidence, billing state, AI metadata, and related product data	account data, workspace content, audit records, billing state, AI runtime metadata, document metadata, integration metadata	Core	Active hosted environment materials reference Neon-backed PostgreSQL topology. Exact region should be confirmed per active deployment.
Cloudflare R2	Object storage	Uploaded files, attachments, previews, and derivative document artifacts	uploaded files, attachments, document artifacts, related object metadata	Feature-dependent	Storage region and transfer path depend on configured environment and deployment topology.
Mailjet	Email delivery	Authentication mail, invitations, notifications, and digests	recipient email addresses, limited email content, template variables, delivery metadata	Feature-dependent	Email processing is limited to workflows that actually send mail. Region and transfer path should be reviewed together with active customer configuration.
Stripe	Billing and subscription operations	Subscription checkout, reconciliation, and billing state updates	billing contact data, provider customer IDs, subscription IDs, commercial metadata, billing events	Commercial-path dependent	Billing follows the enabled commercial path rather than the whole workspace data model.
OpenAI	AI processing	AI-assisted drafting, analysis, conversations, process-building,			

and document interpretation where enabled | prompts, user messages, selected context, assistant outputs, process drafts, document-derived context | Optional / feature-enabled | AI processing is a feature-specific outbound path. Transfer path depends on enabled AI workflow and active provider configuration. |

| Berget | AI processing | Additional AI-assisted service path where configured | prompts, user messages, selected context, outputs, AI run metadata | Optional / feature-enabled | Used only where the relevant AI path is enabled in the environment or feature flow. |

| Google | Customer-enabled integrations | Connected Google workflows such as calendar, mail, and related account-linked automation | external account IDs, email addresses, tokens, scopes, calendar metadata, attendee or event payloads | Customer-enabled integration | Processing path depends on whether the customer enables the integration and what scopes are granted. |

| Pipedream | Managed integration authentication and API proxy | Connects and authenticates customer-selected provider accounts and forwards approved integration requests where the managed path is enabled | organization-scoped connection identifier, connected account identity, provider credentials and scopes, approved request and response payloads | Customer-enabled integration | Used only for integrations configured through the managed connection path. Region and transfer posture depend on the applicable Pipedream service arrangement. |

| GitHub | Customer-enabled integrations | Connected GitHub workflows such as repository-linked operations and account-linked actions | external account IDs, repository metadata, tokens, scopes, commit or repository context where enabled | Customer-enabled integration | Processing path depends on whether the customer enables the integration and what repositories or scopes are connected. |

| Nutrient | Document processing | Document conversion, accessibility analysis, remediation, and related validation workflows where enabled | uploaded documents, extracted or transformed document content, accessibility findings, derivatives and validation metadata | Feature-enabled | Used only by document capabilities configured for the service. Region and transfer posture depend on the applicable service arrangement. |

| Vercel | Application hosting | Hosts application routes and deployment infrastructure | request data, workspace content processed by application routes, operational logs and diagnostics | Core deployment path | Exact runtime region, log location and transfer posture depend on the active deployment. |

| External process worker hosting | Process and document execution | Runs configured long-running process and document jobs | document or workspace payloads required by the job, execution context, callbacks and operational logs | Feature/deployment dependent | Provider, region and transfer posture depend on the active worker deployment. |

| Microsoft | Customer-enabled integrations | Connected Microsoft 365 and Outlook workflows | external account identity, tokens, scopes and selected mail/calendar/contact payloads | Customer-enabled integration | Processing depends on customer authorization and capability-specific scopes. |

| LinkedIn | Customer-enabled integrations | Connected posting or recruitment workflows where provider approval and customer configuration allow them | external account identity, tokens, scopes and selected posting/recruitment payloads | Customer-enabled integration | Processing depends on provider approval, customer authorization and the permitted

integration purpose. |

| WorkOS | Enterprise identity | SSO connection and authentication infrastructure where enabled | identity, organization/domain configuration, SSO connection metadata and authentication records | Optional / organization-enabled | Region and transfer posture depend on the applicable WorkOS service arrangement. |

Interpretation notes

Core database boundary

The default truth path for the product is the application plus its PostgreSQL-compatible database layer. Reviews should start there before evaluating any feature-specific provider path.

Storage boundary

Files and derivative document artifacts can additionally pass to object storage. This is separate from the core relational system of record.

Email boundary

Mail-processing providers receive only the data needed to send the relevant email flow, not the full workspace dataset.

Billing boundary

Billing providers receive commercial and subscription-related data needed to operate the customer relationship, rather than general workspace content.

AI boundary

AI-assisted processing should be reviewed as a scoped feature path. It is not the same thing as saying that the entire application database is routinely sent to an AI provider.

Unless expressly agreed otherwise in writing, Sefira does not use Customer Personal Data to train general-purpose AI models for its own independent purposes.

Integration boundary

Google and GitHub are customer-enabled integrations rather than always-on recipients. Pipedream is used only when the customer selects a managed connection path for an enabled integration. Depending on the data flow, those services may operate under their own terms as customer-selected third-party services, in addition to any processor-side obligations that apply where Sefira sends Customer Personal Data to them on Customer's behalf.

Legal entity and contracting note

This public list identifies the provider brand or service path used in the Sefira environment.

Where a customer requires exact contracting entity detail, enterprise DPA mapping, or deployment-specific confirmation, the relevant legal entity, region, and transfer posture

should be confirmed through the applicable deployment and diligence materials.

International transfers

Some provider paths may involve processing outside the EEA, depending on deployment topology, enabled integrations, and selected AI or storage path.

Where a listed subprocessor processes Customer Personal Data outside the EEA in a country without an adequacy decision, Sefira relies on the applicable EU Standard Contractual Clauses or another valid transfer mechanism under applicable data protection law, together with supplementary measures where appropriate to the service path.

Subprocessor obligations

Sefira's customer-facing DPA requires appropriate written protection for subprocessors. Contract, region, retention and transfer evidence is maintained per active provider path; where diligence is still open, Sefira must not represent the provider as fully verified until a human contract owner records that decision.

Change management and notice

Sefira may update this list as providers change, as deployment topology evolves, or as optional features and integrations are expanded.

For material new subprocessors that process Customer Personal Data on behalf of customers, Sefira intends to provide notice through an update to this page and/or customer communication before the new provider is used, unless a shorter notice period is required for security, continuity, or urgent operational reasons.

Customers with heightened diligence requirements should request customer-specific subprocessor and transfer confirmation where the exact deployment, legal entity, or feature path is material to the review.