Security White Paper

Customer diligence summary

Executive summary

This white paper is the public-facing summary of Sefira’s current security posture.

It is designed to help customers, security reviewers, procurement teams, and privacy stakeholders understand the current control model of the service without pretending that a short public document can replace contract, deployment, or implementation review.

The strongest current security themes in Sefira are:

• privileged administrative access is separately protected;
• sensitive integration secrets are protected with application-layer controls;
• external callback and provider-signature paths are verified rather than blindly trusted;
• the core product truth path is the application plus its PostgreSQL-backed system of record;
• storage, billing, email, AI, and integrations are treated as distinct provider paths rather than one undifferentiated data sink;
• retained-history and anonymization boundaries are handled explicitly instead of being hidden behind unrealistic delete promises.

Service architecture and core data path

For most product workflows, the primary system-of-record path is:

1. an authenticated user or invited collaborator acts inside the application;
2. the application validates and stores core product records in a PostgreSQL-backed database through Prisma;
3. related workflow artifacts, files, or derivative records are written to the storage path configured for that environment; and
4. additional providers are involved only when the relevant feature path requires them, such as email delivery, billing, AI processing, or external integrations.

This matters because customers often want to know whether data stays inside the core application boundary or is passed onward to another processor. In Sefira, the default truth path is the application plus its PostgreSQL-backed database, with additional providers used only for feature-specific flows.

Hosting and provider layers

The active hosted materials for the service describe a Neon-backed PostgreSQL topology for the core relational system-of-record path. The public subprocessor materials additionally identify Cloudflare R2 for object storage, Mailjet for email delivery, Stripe for billing, and configured AI or integration providers where those feature paths are enabled.

This white paper does not claim that every environment is identical. Region, active provider

path, and customer-specific enablement should be confirmed through the applicable deployment, diligence, and subprocessor materials where exact certainty is required.

Environment boundaries

The service includes separate operational concerns for production-facing usage, development change, support handling, and compliance operations. Customer review should treat production data path, provider path, and post-termination handling as distinct questions rather than assuming a single flat infrastructure story.

This public document does not make a blanket claim that all environments are identical or that every operational workflow is customer-visible. Where a customer requires deployment-specific environment detail, that should be handled through scoped diligence follow-up.

Identity, access control, and privileged operations

The service uses authenticated access, role-checked privileged surfaces, and additional portal-gating measures for sensitive administrative operations.

The service applies the following access-control posture:

Control area	Service posture
Authentication	Normal product access requires authenticated access
Privileged admin access	Sensitive portal/admin surfaces are protected by role checks
OTP/MFA-style protection	Portal administrative access is bound to an additional OTP gate linked to the active session
Least privilege posture	Privileged operations are separated from ordinary workspace activity
Auditability	Sensitive governance and anonymization operations produce durable audit/evidence records

This white paper does not claim universal just-in-time access, formal access-review cadence, or independent privileged access certification unless separately stated in customer-specific materials.

The service also applies anti-abuse protections on authentication and privileged access paths. These protections include request throttling, failed-attempt tracking, temporary lockouts that can escalate on repeated abuse, and durable audit evidence for sensitive portal-authentication failures and lockout triggers.

Tenant and workspace boundaries

Sefira is a workspace-oriented SaaS product. The normal security model depends on authorization checks that constrain users to the workspaces, organizations, and roles assigned to them.

Administrative and compliance surfaces are separately protected from normal workspace collaboration flows. Signed storage and callback patterns are also used to reduce the risk

that external paths can bypass normal product trust boundaries.

Secret and token protection

The current implementation includes concrete token-protection measures rather than only policy language.

Publicly describable examples include:

• recoverable integration credentials are protected using application-layer encryption;
• the integration token encryption path uses AES-256-GCM with per-token random IVs;
• sensitive values required for service operation are not treated the same way as ordinary metadata;
• refresh-token and session-binding protections are used for sensitive portal access paths.

At a higher level, the service posture distinguishes between:

• secrets that must remain recoverable to operate an integration path;
• tokens or references that can be hashed or otherwise non-recoverably protected; and
• ordinary product data that should never be treated like a credential.

Encryption and transport protections

The service is designed to protect data in transit and to avoid casual exposure of secrets, tokens, and protected provider credentials.

This white paper describes the following transport and secret-handling posture:

• service and provider communication is expected to rely on encrypted transport;
• sensitive integration credentials are protected with stronger application-layer handling where recoverability is required for feature operation;
• secret values are not intentionally treated as ordinary application content.

More detailed key-management, provider-specific encryption-at-rest, or backup-encryption specifics should be handled through deeper diligence review where the customer needs that level of assurance.

Data flow integrity and external callback verification

The platform uses signed upload grants, signed callbacks, and verification patterns for external-action paths. This reduces the chance that a forged or replayed request can be accepted as a trusted internal event.

The public evidence also supports concrete verification behavior in billing flows, including provider-signature parsing, validation, and duplicate-receipt handling rather than naive trust of inbound webhook payloads.

That distinction matters because a modern SaaS product is not secure if it only protects database writes while treating external provider callbacks as implicitly trustworthy.

Logging, governance, and durable evidence

Sefira’s posture includes governance auditability and durable compliance-operation evidence.

That is important because privacy requests, privileged operations, anonymization work, billing state changes, and incident-oriented follow-up all depend on a traceable history if they are to stand up in customer review or internal investigation.

The service is designed so that:

• privileged and compliance-relevant operations are logged in a durable way;
• anonymization work produces explicit jobs, step-level handling, and retained-history notice;
• product accountability is treated as an operational control, not only a documentation claim.

Vulnerability and change management

This white paper does not claim SOC 2, ISO 27001, or a specific external certification status unless separately stated in customer-specific materials.

What it does describe is a code-and-control posture in which sensitive areas such as integrations, billing callbacks, anonymization, and privileged access are implemented through explicit logic rather than assumed safe by convention.

Customers who require more specific evidence on secure development lifecycle, code review, dependency monitoring, patch prioritization, or independent testing should request that material separately rather than inferring it from this summary.

Incident response and security reporting

Sefira maintains a security posture in which suspicious activity, privileged operations, and security-relevant product events can be investigated through logged evidence and bounded operational controls.

Personal-data incidents are handled in line with the DPA and applicable law, including assessment, containment, investigation, remediation, and customer communication where Customer Personal Data is affected.

Security issues may be reported through the channels identified in the public terms and customer documentation. This white paper does not claim a separate formal bug-bounty or certification-backed incident program unless separately stated.

Backup, recovery, and continuity

Backup, recovery, and service-continuity handling are part of the expected control posture for core persistence layers and provider-backed service operation.

This white paper does not publish formal RPO, RTO, or uptime commitments. Where a customer requires those commitments, they should be addressed in the customer agreement, order form, SLA, or deployment-specific diligence package.

Data classes and retained history

A useful way to understand the security posture is by data class:

Data class	Examples	Current handling theme
Account and identity data	names, emails, membership state	access-controlled product records
Workspace content	tasks, comments, process data, related work history	workspace-scoped authorization and retained-history controls
Documents and derivatives	uploads, previews, extracted text, workflow artifacts	storage-path controls and feature-scoped handling
Secrets and tokens	OAuth credentials, integration tokens, provider secrets	encrypted or otherwise protected, with restricted handling
AI-related data	prompts, outputs, run metadata, selected context	feature-scoped outbound processing rather than universal product export
Audit and compliance data	governance events, anonymization jobs, retained-history notices	durable evidence and accountability controls
Billing and commercial data	subscription IDs, provider events, reconciliation state	limited commercial and operational access path

Sefira’s deletion posture is not marketed as universal instant erasure. Instead, the platform distinguishes among purge, anonymization, tombstoning, and justified retention. This is a stronger and more defensible security and privacy position than promising a simplistic delete model the product cannot actually honor.

AI and integration security posture

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 a model provider.

The service posture is:

- AI-assisted features are assistive and reviewable rather than autonomous authority;
- AI provider path depends on the enabled feature and configured provider;
- customer review is required before AI output is treated as final in material decisions;
- customer and workspace governance should determine which data classes are allowed into AI-assisted workflows.

Integrations such as Google and GitHub are customer-enabled workflow surfaces rather than always-on recipients.

The current public integration security themes include:

- scoped provider authorization is used where supported by the provider;
- integration tokens are protected rather than stored as plain operational values;
- provider-linked workflows are optional and customer-enabled;
- provider callback and webhook paths are verified rather than implicitly trusted.

Assurance status and customer review package

This white paper supports customer diligence. It does not imply a certification or assurance status that is not explicitly stated.

Customers and reviewers who need to understand the service's security posture should review this document together with:

• the Privacy Notice;
• the Data Processing Addendum;
• the Subprocessor List;
• the AI Act Readiness Note; and
• deployment-, provider-, SLA-, or assurance-specific material where those details matter for the review.

Together, those materials make the service reviewable as a real SaaS system rather than as a product described only through high-level security slogans.