AI Act Readiness Note

Customer diligence summary

Executive summary

This note describes Sefira's current AI Act posture for customer-facing review.

It explains:

1. which AI-assisted features are visible in the product today;
2. which AI Act roles are most relevant to Sefira and to customers; and
3. how transparency, provider-path control, reviewability, and change management are handled.

Sefira does not currently provide a general-purpose AI model. Depending on the feature and contractual setup, Sefira may act as a provider or deployer of AI-assisted product functionality, while customers may act as deployers when they use those features in their own business context.

Scope of this note

This note is a public readiness summary rather than a legal opinion or customer-specific classification memo.

It should be read together with the [Privacy Notice](/privacy), the [Data Processing Addendum](/dpa), the [Subprocessor List](/subprocessors), and the [Security White Paper](/compliance/security-white-paper).

The AI Act entered into force on 1 August 2024. According to the European Commission's implementation timeline, the prohibitions, the AI system definition, and AI literacy obligations have applied since 2 February 2025, and the broader AI Act regime becomes applicable on 2 August 2026, including Article 50 transparency obligations.

Current AI feature inventory

Sefira currently uses AI in product areas such as:

• AI workspace conversations and runs;
• draft generation and planning support;
• process-building assistance;
• document-related interpretation, extraction, and summary support; and
• AI runtime review and retrospective analysis support.

The product direction is assistive. AI is used to help draft, analyze, propose, summarize, and structure work. It is not intended to silently redefine domain truth, bypass explicit governance boundaries, or operate as hidden decision authority.

AI Act role assessment

Sefira does not currently present itself as a provider of a general-purpose AI model or foundation model.

Where Sefira uses third-party AI providers to power product features, those upstream providers may have separate obligations under the AI Act that are distinct from Sefira's product-layer obligations.

For Sefira's current product posture, the most relevant role split is:

• Sefira may act as a provider of AI-assisted product functionality where it places those features on the market under its own product and controls the surrounding feature design, routing, and disclosures;
• Sefira may also act as a deployer in its own internal operational contexts, testing, support, or controlled service-operation use; and
• customers may act as deployers when they use Sefira's AI-assisted features in their own business workflows.

That means the important current diligence questions are less about GPAI model-provider obligations and more about intended use, transparency, provider-path clarity, human reviewability, escalation, and customer control boundaries.

Risk classification posture

Based on Sefira's current intended product use cases, the AI-assisted features are not currently designed for the clearest Annex III high-risk AI categories.

That assessment is based on the current visible feature direction: collaboration, drafting, planning, process support, workspace conversations, and document-assisted work.

Current feature-level posture can be summarized as follows:

Feature	Intended use	Likely AI Act category	Current assessment
Draft generation	Drafting, summarization, and planning support	General-purpose assistive use	Not clearly high-risk
Process-building assistance	Workflow support and structured process suggestions	Business productivity support	Not clearly high-risk
Document interpretation	Extraction, summary, and document-assisted understanding	Context-dependent assistive use	Requires customer-context review for sensitive uses
AI workspace conversation	Q&A, drafting, and assistive reasoning	Assistive productivity use	Not clearly high-risk
Retrospective analysis support	Review, interpretation, and operational reflection	Context-dependent assistive use	Requires guardrails where used for significant decisions

Customer-specific use can change the assessment. Separate review is required where AI-assisted features are used in employment, education, essential services, healthcare, legal, financial, public-sector, or similarly regulated decision contexts, especially if the output materially affects individuals.

AI data path and provider control

For AI-assisted workflows, the AI data path is narrower than the full product data path.

The common pattern is:

1. a user interacts with the core application;
2. the application reads the minimum workspace, document, or drafting context needed for the task;
3. that selected context may be sent to the configured AI provider path for the feature; and
4. the resulting output returns to the application and is stored as assistive content, run metadata, or reviewable draft material inside the core product environment.

This means AI processing is not the same thing as the whole application database being exposed to a model provider. The AI path is a feature-specific outbound path layered on top of the core application and its PostgreSQL-backed system of record.

Customers reviewing provider involvement should understand:

• the active AI provider path may vary by deployment, environment, or feature;
• customer configuration may affect whether AI is enabled for a given workflow;
• the relevant providers are identified in the Subprocessor List where they act in a processor-side role;
• customer-enabled integrations remain separate from the AI provider path;
• prompts, selected excerpts, user messages, outputs, and AI run metadata may be processed depending on the feature; and
• unless expressly agreed otherwise in writing, Sefira does not use Customer Content to train general-purpose AI models for its own independent purposes.

Transparency and Article 50 readiness

Article 50 transparency obligations become applicable on 2 August 2026.

Sefira's intended transparency posture is to disclose AI-assisted functionality at the point of use through interface labeling, product documentation, and reviewable AI run context where applicable.

That includes, as relevant to the feature:

• clear labeling that a feature or output is AI-assisted;
• explanatory text near inputs or outputs where a user would reasonably need to understand provider involvement or review expectations;
• reviewable metadata or run context for AI-assisted interactions where the workflow supports it;
• customer-facing documentation describing the feature family and the surrounding controls; and
• disclosure where AI-generated or AI-assisted output is exported, published, or used outside the immediate drafting context, where such labeling is relevant to the workflow.

Sefira's position is that AI-assisted behavior should be understandable to users and

reviewable by customers rather than hidden behind generic automation claims.

AI literacy and internal controls

Sefira maintains role-appropriate AI literacy expectations for personnel involved in designing, configuring, shipping, supporting, or reviewing AI-assisted features.

That literacy posture should cover:

• onboarding to the current AI feature inventory;
• intended-use and risk-classification awareness;
• data-handling and prompt-handling expectations;
• transparency and human-review obligations;
• escalation routes for output quality, misuse, data exposure, or provider-path issues; and
• refresh training when AI features or provider paths materially change.

Human review and reviewability

Sefira treats AI output as assistive rather than authoritative.

The expected governance model is:

• AI-assisted output should remain reviewable before material reliance;
• customers remain responsible for determining when human review is required for their own workflow;
• AI-assisted functionality should not be used as the sole basis for legal or similarly significant decisions about individuals; and
• review context, surrounding workflow state, and auditability should be preserved where the feature design supports it.

Customer responsibilities when using AI-assisted features

Customers remain responsible for how they use AI-assisted functionality in their own business context.

Customers should determine:

• which workflows may use AI-assisted features;
• whether their specific use case could be high-risk or otherwise regulated under the AI Act or sector-specific law;
• what human review, approval, or oversight is required before relying on AI-assisted output;
• whether sensitive or regulated data may be used in the relevant AI workflow under the applicable contract and configuration;
• whether their own users, workers, customers, or counterparties require disclosure; and
• whether additional internal policy, retention, localization, or approval controls are required.

Prohibited and restricted use cases

Sefira's AI-assisted features are not intended for prohibited AI practices, including manipulative, exploitative, social-scoring, or unlawful biometric-categorization use cases.

The product is also not intended to be relied on as a hidden or autonomous decision-maker in regulated contexts without separate review, suitable controls, and an explicit customer-specific assessment.

Monitoring, change management, and escalation

AI Act readiness is not only about the current feature footprint. It also depends on how new or changed AI-assisted features are reviewed over time.

Sefira's intended operating posture includes:

• maintaining an AI feature inventory;
• reviewing new or materially changed AI-assisted features for intended use, provider involvement, data path, transparency, and likely risk classification;
• reassessing customer-facing disclosures when the provider path or feature posture materially changes;
• documenting escalation for misleading output, data leakage, misuse, or provider-path anomalies; and
• routing AI-related issues through support or security handling depending on whether the issue concerns output quality, data exposure, misuse, or provider behavior.

Related materials

Related materials include:

• the Privacy Notice;
• the Data Processing Addendum;
• the Subprocessor List;
• the Security White Paper; and
• any customer-specific or deployment-specific follow-up needed for a particular regulated use case.