

Document reference: EU-04 | Region: EU/EEA | Version: [TBD] | Effective date: [TBD]

WORKING DRAFT - PUBLICATION CONTROL: This document is drafted for the RedTree launch. Do not publish until the Global Compliance Matrix confirms the related launch implementation. Any text marked 'TBD BEFORE PUBLICATION' must be resolved and removed before release.

1. AI transparency

RedTree uses artificial intelligence supplied through the OpenAI API for AI-assisted functions. When a user interacts directly with a RedTree AI feature, the interface will make clear that the user is interacting with or receiving output from an AI system, unless that is already obvious from the context.

2. What the AI may process

To answer effectively, the AI may process relevant user-authorised context such as filenames, folder structure, dates, annotations, document/text content, chat history and user instructions. Rich contextual processing is an intended product function. The data is used for the requested RedTree task, not for behavioural advertising, sale, unrelated profiling or voluntary training/data-sharing programmes.

3. AI output

AI output may be incomplete, inaccurate or misleading. Users should verify important results, especially where decisions may affect health, safety, legal, financial or other significant interests.

4. AI-generated or manipulated content

Where RedTree generates or materially manipulates content and EU AI Act transparency/marketing rules apply to that output, RedTree will provide the disclosures or machine-readable marking required for the relevant use case. TBD BEFORE PUBLICATION: confirm the exact launch set of generative output features and corresponding Article 50 marking obligations.

5. No hidden significant decisions

RedTree does not currently use AI to make hidden decisions with legal or similarly significant effects about users. If a future feature changes that position, it requires a separate compliance review before launch.