REVIEWIT TRUST CENTER · ARTIFICIAL INTELLIGENCE

RITA — AI Transparency Notice

What REVIEWIT's AI assistant does, what it never does, which data it uses, and how you stay in control.

Public document · Version 1.0 — November 2026 · All hosting regions

1. What RITA is

RITA is the assistant built into REVIEWIT. It helps **reviewers, administrators and auditors** understand the governance evidence in front of them — campaigns, access reviews, resources, documentation — and find their way around the product. In **AI-assisted** mode, some answers are written by a generative AI model. In **guided** mode, RITA answers from predefined requests and sourced results, without generative AI.

2. What RITA never does

- RITA never makes, recommends, pre-selects or executes a certification decision — **approve, revoke, escalate or sign**.
- RITA never writes the justification of a decision.
- A **safeguard independent of the model** blocks any answer that would recommend a decision.
- Every decision remains **human**. No automated decision is made about any person whose access is reviewed.
- RITA does not act on your systems: every tool it uses is **read-only**.

3. Data RITA uses

- **Your scope only.** RITA sees only data the authenticated user is already entitled to see, within their workspace. The scope comes from the user's identity, never from the conversation.
- **Minimised data.** In AI-assisted mode, the model receives your question, the context of the page you are on and the aggregated results needed to answer.
- **Masking before sending.** Secrets (keys, tokens, passwords, credentials, payment card and bank account numbers) and special categories of personal data (Articles 9 and 10 GDPR) are masked before anything is sent to the model.
- **No conversation storage.** REVIEWIT does not keep your conversations on its servers. For each exchange, the audit trail records technical information (model, tools consulted, volume) and a fingerprint of your question — not its text.
- **No training.** REVIEWIT does not use your data to train or fine-tune any model. **The models are served by our hosting provider's model service, whose contractual terms exclude any use of your content to train or improve models; the model provider has no access to your requests or answers and receives none of your content.**

4. Where processing takes place

REVIEWIT serves customers worldwide from hosting regions in the **European Union, the United States, India, Japan and Australia**. AI models are served within the geography of your subscription's hosting region, and a hosting region only calls models from its own geography: **a subscription hosted in Europe has its data processed in Europe only.** In guided mode, no data leaves REVIEWIT. Model providers and our hosting provider are listed in the list of sub-processors published on the Trust Center.

5. You stay in control

- The **subscription administrator chooses the mode**: Disabled, Guided (no AI) or AI-assisted. A workspace may restrict this choice, never exceed it.
- AI-assisted mode requires acceptance of an **AI data processing addendum**.
- **Kill switch**: the administrator can switch RITA off at any time; it takes effect within a minute.
- REVIEWIT can switch off a model family across the platform **in under a minute**; the reason is displayed to customers and RITA continues in guided mode.

6. Quality and limits

Every model route (model, version, geography) is evaluated against security, accuracy and fairness thresholds **before it is served**, and again whenever RITA or the model changes; a model that does not meet the thresholds is

not served. Methods and results are set out in the **RITA AI Safety & Robustness Evaluation Summary**. An answer written by AI may still be inaccurate or incomplete: check important points against the sources shown under the answer.

7. Regulatory position — EU AI Act

REVIEWIT places RITA on the market under its own name and is the **provider** of this AI system; customers who use RITA under their own authority are **deployers**. REVIEWIT integrates general-purpose AI models without modifying them substantially and is therefore not a provider of a general-purpose AI model. RITA engages in **no practice prohibited by Article 5** (no manipulation, no social scoring, no emotion recognition, no biometric categorisation). RITA is used in the context of employee access management (Annex III, point 4): it neither takes nor recommends decisions and provides procedural assistance only. Risk classification: **not high-risk under Article 6(3): RITA performs a narrow procedural task, involves no profiling and does not materially influence the outcome of decisions, which remain human. This assessment is documented (Article 6(4)) and RITA will be registered in the EU database (Article 49(2)). Obligations applied: transparency (Article 50) and AI literacy (Article 4).**

8. Report a problem

If an answer seems inaccurate, inappropriate, or appears to go beyond your scope, tell us at **dpo@we-reviewit.com**. Every report is reviewed on receipt. If needed, RITA is suspended for the subscription or model concerned while the issue is analysed, and is re-enabled only after a new evaluation.