CORTEX · SUPPORT INTELLIGENCE ENGINE

Cortex and the EU AI Act: *not high-risk, by design.*

Ascentis AI has assessed Cortex against the EU AI Act, Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744. This statement sets out the outcome, the obligations that apply, and how Cortex meets them.

SELF-HOSTED AI · OPEN-WEIGHT MODELS · A PERSON SENDS EVERY AI DRAFT

Not high-risk	None	Art. 4 & 50	0
ARTICLE 6 CLASSIFICATION	PROHIBITED PRACTICES ARTICLE 5	THE OBLIGATIONS THAT APPLY — MET	EXTERNAL AI CALLS

Self-assessed by Ascentis AI against the text of the regulation · issued 4 November 2026

THE STATEMENT

"Ascentis has assessed Cortex as not high-risk under Article 6 of the EU AI Act and meets its Article 4 and Article 50 obligations. All AI runs on open-weight models on the server Cortex is installed on; nothing is sent to a cloud AI service. No AI-written email reaches a customer until a person reviews and sends it, and every automatic AI effect is reversible and audited."

HOW THE CLASSIFICATION WAS REACHED

Tested against every route to *high-risk*.

01

Prohibited practices — none

Cortex does not manipulate or exploit anyone, score people socially or predict crime. It holds no biometric data. It reads sentiment from the **text** of a survey reply, which is not emotion recognition as the Act defines it.

02

Product safety (Annex I) — not met

Cortex is not a product covered by EU product-safety legislation, nor the safety component of one. It is a support-desk system that suggests, drafts and records. It controls no machine.

03

Listed uses (Annex III) — none

Each AI use was checked against the areas it could touch: essential services, creditworthiness, allocating work and evaluating workers. None is listed, so no exemption is relied on and no registration applies.

ANNEX III · USE BY USE

Every AI use, checked against the areas it *could touch*.

AI USE IN CORTEX	AREA CONSIDERED	WHY IT IS NOT HIGH-RISK
Reading a case — product, customer, details and a one-line summary	–	Text extraction for a support ticket. No listed area applies.
Suggested priority, which sets the SLA target	5 · Essential services	Business-to-business support for equipment, not access to essential services for individuals. A person can change it.
Root cause, resolution suggestions, reply drafts, Ask Cortex	–	Decision support for engineers. Every email reaches a customer only through a person.
Pattern alerts, spam hiding	–	Case housekeeping. Each effect is reversible on the case.
Commercial verdict — billable or covered	5(b) · Creditworthiness	A rule engine over the contract, warranty and hours; the AI supplies only an editable hours estimate. It concerns business billing, not individuals' credit.
Who gets a new case	4(b) · Allocating work	Decided by your routing rules, availability and capacity — not by the AI, and not on behaviour or personal traits.
Satisfaction ratings in points, badges, the leaderboard and performance views	4(b) · Evaluating workers	**Designed out.** A rating counts toward a person's scores only when the customer typed it, or a colleague — never the rated engineer — confirmed it. An unconfirmed AI reading scores no one.
Speech to text, OCR	–	Transcription of the user's own input.

CONCLUSION

Cortex is *not a high-risk AI system*. Its obligations are AI literacy under Article 4 and transparency under Article 50 — and it is re-assessed before any new AI use ships.

TRANSPARENCY · ARTICLE 50

AI-written email is *marked and reviewed*.

- **Ask Cortex says what it is.** The assistant's panel states that its answers are AI-generated from your Cortex data and manuals. ART 50(1)
- **AI-written email is marked in the message.** Every draft records whether the AI, a template or a person wrote it. An email sent from an AI-written draft carries the machine-readable header `X-AI-Disclosure: mode=ai-originated` over SMTP and Microsoft 365. It is invisible to the reader and names no model. ART 50(2)
- **AI text on screen is labelled.** Summaries, suggestions, alerts and the staff digest stay with your team and are marked as AI.
- **Automatic emails carry no AI text.** The chase, clarification, satisfaction survey and visit emails are fixed templates. The chase and clarification quote only the subject line of the customer's own first email.

AI LITERACY · ARTICLE 4

Your team knows *what the AI does*.

- **For your people.** Every deployment ships with the *AI in Cortex* manual: each AI feature, what it can get wrong and how to override it, what runs automatically and how to undo it, measured accuracy with its date, and where the AI runs. Ask Cortex answers from it.
- **For ours.** Everyone at Ascentis who builds, sells, installs or supports Cortex reads the same manual and the full compliance file.

NOT APPLICABLE

Emotion recognition and biometric categorisation ART 50(3), and deep fakes or text published to inform the public ART 50(4). Cortex does neither.

MODEL REGISTER · AS AT 4 NOVEMBER 2026

Open-weight models, *on the Cortex server*.

MODEL	WHAT CORTEX USES IT FOR	LICENCE
Qwen3.5 35B · Alibaba	Every language task: reading cases, priority, root cause and resolution suggestions, drafts, Ask Cortex, reading survey replies	Apache-2.0
EmbeddingGemma · Google	Search embeddings	Gemma Terms of Use
BGE reranker v2 m3 · BAAI	Re-ranking search results (optional)	Apache-2.0
Whisper small · OpenAI	Voice notes to text	MIT
Tesseract	Text from scanned documents and images	Apache-2.0

Every model runs on the server Cortex is installed on. No inference, embedding or transcription is sent to a cloud service, and no customer data trains any model.

ROLES UNDER THE ACT

Who is *responsible for what*.

- **Ascentis AI is the provider** of Cortex, placed on the market under its name.
- **Model providers keep their obligations.** Cortex uses open-weight general-purpose models unmodified, with no fine-tuning; Articles 53–55 sit with their providers.
- **You are the deployer.** The manual sets out your duties: AI literacy for your staff, your data protection impact assessment, consulting staff if you use the leaderboard, and re-measuring if you edit prompts or models.

CHANGE CONTROL

The assessment moves *with the product*.

- **A new model** is re-measured on the 120-case benchmark before it ships. This register and the *AI in Cortex* manual are updated in the same change, with the licence read from the model card.
- **A new AI use** is classified against Annex III before release, not after. Prompts are versioned, and the version is logged with every classified question.
- **Your own changes:** if you edit prompts or switch models, Ascentis's measured figures no longer apply. The manual tells you so.

ISSUED BY

Erwan Lhermitte

Erwan Lhermitte
Founder & CEO, Ascentis AI Ltd

4 NOVEMBER 2026

A STATEMENT, NOT A CERTIFICATE

This is Ascentis AI's own assessment against the text of the regulation. The EU AI Act provides conformity assessment only for high-risk systems, so there is no certificate to hold for a system like Cortex — and we do not claim one.

Questions on this statement, or on *your duties as a deployer*?

INFO@ASCENTIS-AI.COM · ASCENTIS-AI.COM