

ORamaVR SA

Responsible AI Policy

Human oversight, transparency and accountable use of AI in ORamaVR products

Document ID: OVR-POL-013 · v1.2 · 5 August 2026

1. Status, purpose and scope

1.1 This Responsible AI Policy (OVR-POL-013) is an ORamaVR Standard-Form Document, published at oramavr.com/legal/responsible-ai. It states the principles ORamaVR applies to the artificial-intelligence features of its products and is referenced by the Code of Conduct & Ethics (OVR-COR-001). It is read together with the JARIA Terms of Use (OVR-POL-005), the Acceptable Use Policy (OVR-POL-004), the Fair-Use limits in the Acceptable Use Policy (OVR-POL-004, Annex FU), the Sub-Processor Register (OVR-POL-008) and the JARIA Addendum (Annex J to the JARIA Terms of Use, OVR-POL-005). Capitalised terms not defined here have the meanings given in the ORamaVR Master Glossary (OVR-GLO-001).

1.2 Scope. This Policy covers the AI features of ORamaVR products: JARIA (the in-session intelligent assistant on the VTC Platform, in Base and JARIA+ configurations), OMEN (the generative authoring co-pilot in ORamaVR Creator, for text/content authoring and image generation), and any other AI-assisted features of Creator products. It applies to ORamaVR and, where incorporated into an agreement, to customers, partners and users of those features.

1.3 ORamaVR may update this Policy by publishing a revised version at the URL above; a change that materially reduces a protection stated here takes effect on not less than thirty (30) days' notice, and any other change on not less than ten (10) days' notice, in accordance with the Standard-Form Document Amendments mechanism in the applicable Principal Agreement. Notices: legal@oramavr.com.

2. Human oversight

2.1 AI assists training; humans decide. ORamaVR AI features support education and training. JARIA provides in-session guidance, explanation and feedback; it does not autonomously evaluate, grade, certify or sanction any trainee's clinical competence. Any decision that produces a legal or similarly significant effect on a trainee — assessment outcomes, certification, progression, disciplinary action — is made by the responsible human instructor or institution, not by the AI (see OVR-POL-005).

2.2 Competency reports and analytics generated on the VTC Platform are decision-support inputs for qualified human reviewers. Customers must maintain meaningful human review wherever AI outputs inform decisions about individuals.

3. Architecture and sub-processing

3.1 Managed inference. JARIA in-session LLM inference and OMEN image generation are provided through managed third-party model providers engaged by ORamaVR as sub-processors. The canonical, current list of sub-processors — including providers, roles, locations and transfer safeguards — is the Sub-Processor Register (OVR-POL-008), published at oramavr.com/legal/sub-processors.

3.2 BYOK for OMEN text. OMEN text and content authoring operates on a bring-your-own-key (BYOK) architecture: prompts are routed through the customer's own third-party account and API key, under the customer's own agreement with that provider. That provider acts for the customer, is not an ORamaVR sub-processor, and is not listed as one in OVR-POL-008 (see OVR-CRE-001 §5.5).

4. Training data

4.1 Customer content — prompts, uploaded corpora (including IFU and reference documents), session data and outputs — is not used to train or fine-tune underlying AI models. This is the default-off position stated in the JARIA Terms of Use (OVR-POL-005 §5.3). Any use of customer content for model training requires the customer's express prior opt-in recorded in a separate written agreement; the fine-tuning permission mechanism in the VTC End-Customer Terms (OVR-SCH-002) operates only through that opt-in path. ORamaVR may process aggregated, de-identified usage telemetry for operational, security and general product-improvement purposes in accordance with the Privacy Policy (OVR-POL-002).

5. Transparency

5.1 Users are informed when they are interacting with an AI system: JARIA identifies itself as an AI assistant, and AI-generated content produced by OMEN or other generative features is labelled as AI-generated within the authoring environment. Outputs of generative features must not be presented as authored or clinically validated by a qualified human professional where they have not been (OVR-POL-005 §6).

6. EU AI Act positioning

6.1 ORamaVR products are XR simulation tools for the education and vocational training of healthcare professionals. They are not medical devices (see the Regulatory Positioning Memo, OVR-REG-001) and do not perform autonomous high-risk decision-making concerning individuals: assessment and certification decisions remain with human instructors and institutions (Section 2). On this basis ORamaVR's working position is that its AI features are not high-risk AI systems under Regulation (EU) 2024/1689 (the EU AI Act) Annex III.

6.2 ORamaVR nonetheless maintains internal technical documentation for JARIA structured on the Annex IV headings (OVR-REF-AI-001), and complies with the transparency obligations applicable to AI systems that interact with natural persons and generate content. Classification is reassessed on any material feature change or regulatory guidance affecting it.

7. Quantitative limits

7.1 Quantitative usage caps on AI features (including OMEN image-generation volumes and Creator Solo interaction limits) are set exclusively in Annex FU (Fair-Use Limits) to the Acceptable Use Policy (OVR-POL-004), which is the single home for such caps and prevails over any figure in product collateral.

8. Incident reporting

8.1 Suspected malfunction, misuse or harmful output of an AI feature — including output that could mislead a user on a clinically significant point — should be reported to security@oramavr.com (security incidents) or legal@oramavr.com (other concerns). ORamaVR triages AI-related incidents under its incident-response process, corrects or withdraws affected features where warranted, and makes any notifications required by applicable law, including serious-incident reporting where the EU AI Act so requires.

9. Not a medical device

9.1 No ORamaVR AI feature provides medical advice, diagnosis or treatment, and none is intended for use in the diagnosis, treatment, prevention or mitigation of disease. The regulatory positioning of ORamaVR products, and customer responsibilities where customer content changes an intended purpose, are set out in the Regulatory Positioning Memo (OVR-REG-001). No patient-identifiable data is to be processed on the platform (OVR-SCH-002 §4).

10. Ownership and review

10.1 This Policy is owned by ORamaVR Legal/Compliance with the CTO organisation and is reviewed at least annually, on any material AI-feature change, and on relevant regulatory developments. It is governed by Swiss law. Questions: legal@oramavr.com.