

Deploy AI in a public service, and be able to account for it answer by answer.


ELYSÉA fits into the administration's AI tool: a check before and after the model, prohibitions that no publisher can lift, and an attestation on every answer that inspectors can read back.

YOUR PROBLEM

- **Obligations specific to public authorities.** For a high-risk system, registration comes before use; otherwise they “n'utilisent pas ce système” (do not use that system). High-risk systems already in service must comply by 2030-08-02 at the latest.^{1,2}
- **A fundamental rights impact assessment** for public bodies deploying a high-risk system; it may now refer to the data protection impact assessment.²
- **Transparency towards users**, due since August 2026: they must know they are dealing with an AI; in Switzerland, the person “a le droit de savoir s'il parle ou écrit à une machine” (has the right to know whether they are talking or writing to a machine).^{3,4}
- **Protecting staff:** inferring emotions at work has been prohibited by the AI Act since 2025-02-02.¹

HOW ELYSÉA ANSWERS IT

The user talks to the application, never to ELYSÉA



The administration's tool receives the message; ELYSÉA checks before and after the model; every answer goes out with a verifiable attestation.

YOUR REQUIREMENT	WHAT ELYSÉA PROVIDES
Classify the system	ELYSÉA is not a high-risk system: no biometrics, no emergency call triage, no care purpose. Your final use remains for you to classify.
Account to an inspector	A per-answer attestation, tied to the software version; an audit log kept for six months, then purged.
A policy the publisher cannot bypass	Rules can be tightened, never weakened: any attempt is refused.
Protect staff	Emotions locked within the layer (a voice analysis path remains on the deployed service, removed with voice in the upcoming version); referral to public numbers in case of distress.
The model	Today, ELYSÉA's model, which the upcoming version calls in Paris; on the deployed service, the provider setting is not verified; direct access to the model, possibly outside the EU, remains open by configuration; using your own model key (Mistral, OpenAI, Gemini or Claude) is planned, not yet enabled.

WHAT YOU CAN VERIFY TODAY

- PROVIDED ON REQUEST, BEFORE ANY CONSULTATION
- **ELYSÉA's impact assessment and record of processing:** an authorisation under conditions, with thirty measures required before the first real user.
 - **The list of processors**, with each one's region: proven for five, not proven for the others, shown as such.
 - **The measurement protocol**, written and sealed on 2026-08-28 before any measurement, amended since, to be resealed before the campaign; and, in public, the comparator demonstration at elysea.app/comparateur (illustrative data, labelled as such).

What is missing, stated plainly: no measured result (published after the campaign), no third-party validation, processing contracts not signed, entity not registered. The path to a third party: a public benchmark, decided and to be built, that an independent laboratory can re-annotate. Next proofs, in order: upcoming version deployed and verified, pilot, new sealing and publication of the protocol, campaign, results published and re-annotatable by an independent laboratory in a public benchmark still to be built, entity registered.

WHAT SETS ELYSÉA APART

EXISTING APPROACH	WHAT IT DOES, ACCORDING TO ITS SOURCE	WHERE ELYSÉA DIFFERS
Governance platforms: Credo AI, Holistic AI	Inventory, regulatory mapping (AI Act, ISO 42001), supervision.	They document compliance; ELYSÉA produces the proof in every answer. Complementary. ^{5,6}
A cloud provider's content filters: Bedrock Guardrails	Filtering policies configurable by the team.	Conduct prohibitions that cannot be switched off, a per-answer attestation. ⁷
Rules developed in-house: NeMo Guardrails	An open framework to program your own checks.	A base already written, which the publisher cannot lower. ^{8,9}

OFFER AND COMPLIANCE

Offer

- **Equipping staff:** Teams offer, €30 excl. VAT per seat per year up to 5,000 seats, then €18 and €9.60, by volume.
- **Serving users** through a publisher: builder offer, per active user.
- **Public procurement:** no bid before the entity is registered.

Compliance

- **Staff representative body:** consulted before deployment, as a company would.
- **Tender specifications:** require the provider's impact assessment, its record of processing and its list of processors with their regions.
- **Data:** database in Paris, execution in Belgium, model in Paris, accounts in Frankfurt; the upcoming version removes from the core several paths to services possibly outside the EU, other providers remain with an unproven region. ELYSÉA does not call itself sovereign.

NEXT STEP

A 60-minute joint reading with your data protection officer and your IT department: impact assessment, record of processing, processors. Then a pilot with a publisher, using a test key. contact@elysea.app · elysea.app

SOURCES (web pages read on 2026-10-01 unless dated) 1 EU — Regulation 2024/1689, OJ L of 2024-07-12 2 EU — Regulation 2026/1744, OJ L of 2026-07-24 3 European Commission, updated 2026-08-03 — AI Act — regulatory framework 4 PFPDT, 2025-05-08 — Update : la loi actuelle est applicable à l'IA 5 Credo AI — Home page 6 Holistic AI — Home page 7 AWS — Amazon Bedrock Guardrails 8 NVIDIA — NeMo Guardrails — documentation overview 9 NVIDIA, version 0.24.1, accessed 2026-10-01 — NeMo Guardrails — public repository