Sovereign AI for regulated enterprises: GDPR and the EU AI Act

20 May 2026

Regulated organizations want the benefits of modern AI, but the usual route sends personal and confidential data to a cloud model run by a third party. Under GDPR and the EU AI Act, that route adds legal weight to every deployment. Running AI on your own infrastructure removes most of that weight at the source. This briefing sets out how.

The problem with the cloud route

When prompts and documents leave your network to reach a hosted model, the model provider becomes another party to your data. That raises questions a regulated buyer has to answer: what is the lawful basis, who is the processor, where does the data sit, and how is any of it auditable. Each question is answerable, but each adds contracts, assessments, and risk. Sovereign, on-premise AI answers most of them by never letting the data leave in the first place.

GDPR

GDPR is easier to satisfy when personal data stays inside the controller's own boundary.

- No third-party processor for the model. Because the model runs on your infrastructure, there is no external processor handling personal data on your behalf for inference, and no processing agreement or onward-transfer chain to manage for it.
- No international transfers. Data that never leaves your premises is not transferred to another country, so the transfer rules that complicate cloud AI do not apply.
- Data minimisation and residency by construction. The retrieval index and the logs live where you put them, in your jurisdiction, under your retention rules.

- A cleaner lawful basis. With no secondary use of the data, the basis for processing is the purpose you already have, not a new one created by sending data to a vendor.

The EU AI Act

The AI Act places duties around transparency, governance, record-keeping, and human oversight, scaled to the risk of the system. On-premise deployment does not by itself classify or certify a system, but it makes those duties easier to meet.

- Transparency and records. Because the models, the index, and the logs are inside your boundary, you can see and keep the records the Act expects, rather than relying on a provider's summaries.
- Human oversight. The system runs where your people and controls are, so oversight is a matter of your own processes rather than a vendor's console.
- Governance. One boundary to govern is simpler than a set of external integrations, which makes the assessment and monitoring the Act calls for more tractable.

Confidentiality and sector rules

Beyond the two regulations, regulated work carries client NDAs and sector rules, in finance, healthcare, the public sector, and legal services. Material processed in place stays inside the confidentiality you already promised, and does not become a question of a vendor's controls.

Verifiable, not asserted

Compliance is a legal judgement made in your context, and this briefing is not that judgement. What Empisto provides is an architecture that makes the judgement easier and the evidence available: the platform runs inside your walls, has no outbound path for your data, and is open to independent audit of its data handling and networking. Trust is meant to rest on that evidence rather than on a vendor's word.