Private AI for regulated industries: on-premise LLMs for law, health and finance

26 October 2026 · 6 min read · blogs.iictp.org/private-ai-on-premise-llms-regulated-industries

Law firms, clinics and financial teams handle material they may not be free to send to an outside service. Running a model on infrastructure you control can keep that material inside your boundary, but it does not make a deployment compliant by itself.

TL;DR The short version

- Running a model on your own infrastructure can keep sensitive legal, health or financial material inside your boundary.
- It does not make a deployment compliant by itself: you still need access control, logging and retention rules.
- Answer first what data the system touches, under which rules, and who may use it for what; regulations differ by country and sector.

Read online, listen to it, or change the text size: https://blogs.iictp.org/private-ai-on-premise-llms-regulated-industries

Why organisations choose it

- Data stays on systems the organisation controls.
- Data-residency and contract limits are easier to meet.
- The system can run in a network with no outside access.
- Costs are predictable once the hardware exists.

What you still have to do

- **Access control.** Restrict who can use the model and what documents it can read.
- **Logging and retention.** Decide what is recorded, who sees it and for how long.
- **Isolation.** Segment the network; log and review administrator access.
- **Updates.** You own patching for the runtime, the drivers and the model.

Questions to answer first

1. What data will the system touch, and under which rules?
2. Who may use it, and for what?
3. What happens when it is wrong, and who checks?

4. Who is accountable for the system?

A note on compliance

Regulations differ by country and sector. Ask your compliance or legal team what applies. Nothing in this post is legal advice.

> Related credential: Local LLM Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. NIST. AI Risk Management Framework (AI RMF 1.0). nist.gov/itl/ai-risk-management-framework.

Published by IICTP, an independent technical credential body, on 26 October 2026. The current version of this post is at https://blogs.iictp.org/private-ai-on-premise-llms-regulated-industries. More guides: blogs.iictp.org.