



AI agents leaking your data.

The hidden risks of public LLMs, and
how Private Agentic Flows are the
solution.



Rodrigo Rubio
Head of Digital



The Problem with Public AI

Agentic AI can reason and take complex actions, but if you operate in Healthcare, Finance, Legal, or Defense, consumer-facing LLMs and public APIs are a massive dealbreaker due to privacy and compliance concerns.

Sending protected information to third-party models risks violating frameworks like **HIPAA** and **GDPR**.



The Solution

To harness the power of AI without compromising security, you need agents that operate entirely behind your firewall.

This keeps your sensitive proprietary data secure while still allowing the AI to reason and execute complex tasks on your behalf.



How is it built?



The 3-layer-system

- 1** Foundation Layer: Your LLM running entirely on your own infrastructure (on-prem or private cloud).
- 2** Augmentation Layer: Your RAG, Vector Databases, or fine-tuned adapters that retrieve context from your private knowledge bases.
- 3** Action Layer: The secure tools and internal APIs your agent uses to execute functions and interact with your databases.



What to look out for



Embedded data in fine-tuning

The Threat: When you fine-tune models on private data, that information gets embedded into the model. If accessed, it can be extracted.

The Fix: Data Anonymisation. Scrub all Personally Identifiable Information (PII) before it ever touches your LLM. Replace names with tokens or hashes to ensure data is anonymous from the get-go.



Too much access

The Threat: Giving an AI agent access to an entire database when it only needs one data point.

The Fix: Data Minimisation. Only give your agents access to the absolute minimum data required for their specific task. For example, a scheduling agent doesn't need a full medical history, just names and availability



The Future of Enterprise AI is Private.

Are You Ready?



As AI becomes integrated into critical workflows, the question isn't if you should build private agentic flows, but how quickly you can get there.

Navigating the complexities of secure LLM infrastructure, data anonymisation, and compliance isn't easy. Consult tedix today to design and execute the right implementation strategy for your secure, private Agentic AI ecosystem.



Let's build AI that works for you, and only you.

Start your secure AI implementation with Tedix.

Get in touch!



Rodrigo Rubio