The AI Reality: Why Renting Your Intelligence is The Biggest Risk

Back in September, I asked a critical question: **The £50 Million Question: Why Your AI Strategy Needs to Stay In-House.**

Three months later, the evidence is overwhelming. The answer is clearer than ever: **security, control, and measurable ROI are not optional, they are existential.**

The Hype Is Crumbling.

The ROI Question Nobody Wants to Answer

The AI chatbot race is "really heating up", Mistral just unveiled powerful new models, and OpenAI's Sam Altman declared a "code red", admitting rivals are closing in on ChatGPT's early lead (*FT News Briefing, 3 December 2025*).

But here's what matters more: Palantir CEO Alex Karp delivered a stark warning that many AI projects "may not create enough value" to justify their massive costs (*Yahoo Finance Invest, November 2025*). He distinguishes between basic applications that fail to move profit margins and a results-driven subset that delivers measurable business value.

The real question is not who has the fastest model. It's **whether these models are useful and will get adopted** (*FT News Briefing*).

Copyright © 2025 Wenceslao Alfageme/AlfaFinTec. All rights reserved.

The Scaling Illusion is Fading

The "scaling is all you need" optimism that fuelled the AI boom is hitting a wall. In a survey of 475 AI researchers by the Association for the Advancement of Artificial Intelligence, **76 per cent** said it's "unlikely" that simply expanding current models will achieve Artificial General Intelligence (AGI) (*New Scientist, 14 March 2025*).

Stuart Russell at UC Berkeley put it bluntly: "The benefits of scaling in the conventional sense had plateaued".

And yet, tech companies plan to collectively spend an estimated **$1 trillion** on data centres and chips in the next few years.

Perception vs. Reality

80 per cent of survey respondents believe current perceptions of AI capabilities "don't match reality". As Thomas Dietterich at Oregon State University noted: "Systems proclaimed to be matching human performance still make bone-headed mistakes. These systems can be very useful as tools for assisting in research and coding, but they are not going to replace any human workers" (*New Scientist*).

The Security Crisis is Structural, Not Incidental

If the ROI concerns weren't enough, the security landscape should keep every executive awake at night.

Critical security vulnerabilities have been uncovered in AI inference engines from **Meta, Nvidia, and Microsoft** (*Oligo Security, November 2025*). The pattern, called "ShadowMQ", could allow attackers to execute arbitrary code on AI inference servers.

The worst part? **Widespread code reuse** propagated these flaws across the AI ecosystem. As researcher Avi Lumelsky warned: "When code reuse includes unsafe patterns, the consequences ripple outward fast".

Oligo identified **thousands of exposed connections** on the public internet linked to vulnerable inference clusters. Every cloud AI query you send is potentially a liability.

The Solution: Own Your AI Strategy

The only way to navigate these risks and unlock genuine, results-driven value is an **In-House AI strategy** using **Local LLMs** that provide absolute data sovereignty.

Here's the pragmatic approach I advocated in September, and the one I'm building:

1. Local LLMs for Sensitive Data Handle analysis of your confidential, strategic data entirely in-house. No processed data leaves your infrastructure. No queries are logged externally. No foreign entity can subpoena your AI interactions.

2. External LLMs via API for Non-Sensitive Tasks Continue using external LLMs strategically, but only for processing or finding non-sensitive data subcomponents or publicly available information.

3. Leverage Efficient Open-Source Models In-house deployment is more feasible than ever. New Chinese models demonstrates an **efficient Mixture-of-Experts (MoE) architecture**, delivering competitive performance at roughly **8% of the cost** of proprietary alternatives like Claude Sonnet 4.5 (*Artificial Analysis, October 2025*).

Copyright © 2025 Wenceslao Alfageme/AlfaFinTec. All rights reserved.

The Commoditisation of General Knowledge

One final wake-up call: Your competitive advantage lies in what you uniquely know. Don't rent it out to systems beyond your control.

Your Next Move

If you are ready to move beyond the hype and prevent your intellectual property from becoming part of a data ecosystem you no longer control:

Commission an AI usage audit → Identify your highest-risk AI interactions → Run a pilot local LLM project with clear success metrics

The question is not whether to make this shift. It's whether you'll lead the change or be forced into it by the next headline-grabbing breach.

Sources Referenced:

1. **FT News Briefing** (3 December 2025) - AI correspondent Melissa Hecula on the AI chatbot race, Mistral's announcement, and Sam Altman's "code red"
2. **Palantir CEO Alex Karp at Yahoo Finance Invest** (13 November 2025) - Warning on AI investment ROI
3. **New Scientist** (14 March 2025) - Survey of 475 AI researchers on AGI scaling skepticism (AAAI report)
4. **Oligo Security Disclosure** (15 November 2025) - ShadowMQ vulnerabilities in Meta, Nvidia, Microsoft AI frameworks
5. **Artificial Analysis / MiniMax M2 Release** (28 October 2025) - Open-source MoE model performance and efficiency
6. **Fortune / Axios** (November 2025) - Karp on domain expertise vs. general knowledge

#EnterpriseAI #LocalLLM #AIStrategy #DataSovereignty #Cybersecurity #DigitalTransformation #Innovation

Disclaimer and Disclosure

Third-party Content and AI Assistance: This article references tools and software that are publicly available and proprietary to their respective creators. It also references publicly available research and news. The author does not claim ownership or affiliation with these third-party products. This article was written by the author with assistance from Generative AI Language Models.

Transparency Notice: While every effort has been made to ensure accuracy, readers should verify information independently and consult official sources or documentation for the mentioned tools and software. The use of AI in the writing process is disclosed in the interest of transparency, but all opinions and analyses are the author's own unless otherwise stated.

Copyright © 2025 Wenceslao Alfageme/AlfaFinTec. All rights reserved.