ISSUE 01 — SPRING 2026

Private AI Europe

The AI sovereignty issue

Why European businesses must own their AI — before regulation decides for them.

ABOUT THIS MAGAZINE

Private AI Europe

Private AI Europe is a publication by **Human Technology eXcellence**. We help European businesses deploy AI that is private, compliant, and under their control. Analysis, case studies, and regulatory insights for decision-makers navigating the AI transition.

privatechatgpt.eu | ht-x.com

IN THIS ISSUE

EDITORIAL

Private AI is not a luxury — it's infrastructure

When I founded Human Technology eXcellence (HT-X), I had a conviction that many considered premature: *the real competitive advantage in AI would not come from the models themselves, but from the ability to deploy them under full control.*

Today, that conviction has become regulation. The AI Act is law. The GDPR enforcement on AI tools is intensifying. And yet, the vast majority of European businesses still send their most sensitive data — client records, financial projections, proprietary processes — to American servers, through tools they neither own nor control.

The numbers are stark. According to the AI Observatory at Politecnico di Milano, the Italian AI market reached €1.8 billion in 2025, growing at +50% year-over-year. But 75% of small businesses haven't even started. Meanwhile, 47% of workers already use AI tools — and 81% of them use personal accounts, outside any corporate governance. *Shadow AI* is not a future risk. It is happening now, in every office, with every prompt.

This magazine was born from a simple idea: European businesses deserve to use artificial intelligence *without* becoming regulatory experts, *without* sending their data across the Atlantic, and *without* depending on a single vendor's pricing decisions.

At Human Technology eXcellence, we build the tools that make this possible, but we are not alone. On-premise AI platforms, running open-source models — Llama, Mistral, DeepSeek, Qwen — directly on the client's infrastructure is a reality. No data leaves. No third party trains on your prompts. Compliance is not a checklist you manage — it's a feature of the system.

In this first issue of *Private AI Europe*, you will find what we believe every business leader needs: real case studies (not theory), regulatory analysis (not panic), and a technology roadmap that puts European businesses in the driver's seat.

The AI race is not about who adopts first. It's about who adopts *well* — with sovereignty, security, and strategic clarity.

I hope this magazine helps you do exactly that.

Francesco Menegoni

CEO, Human Technology eXcellence
Innovation Expert, Italian Presidency of the Council of Ministers
PhD in Engineering, Politecnico di Milano

MARKET

AI Market Europe 2025: Italy, France, Germany, Austria and Switzerland between growth and incomplete adoption

Data from Politecnico di Milano, Bitkom, Bpifrance and Anthropic reveal the same gap across Europe: AI capabilities far outpace actual business usage.

By Francesca Spazzali, MSc

Italy

The Artificial Intelligence Observatory at Politecnico di Milano released updated data in February 2026. The numbers confirm strong growth — but also a widening gap.

€1.8 billion: the value of the Italian AI market in 2025, growing +50% versus 2024. The three-year CAGR 2022–2025 is +54%. Generative AI is the main driver, but not the only one: 54% of the market still comes from traditional Machine Learning projects.

The most relevant finding for businesses: *77% of the market involves custom projects*, built bespoke for individual clients.

Italian AI market by solution type (2025)
Source: Politecnico di Milano — Osservatori Digital Innovation

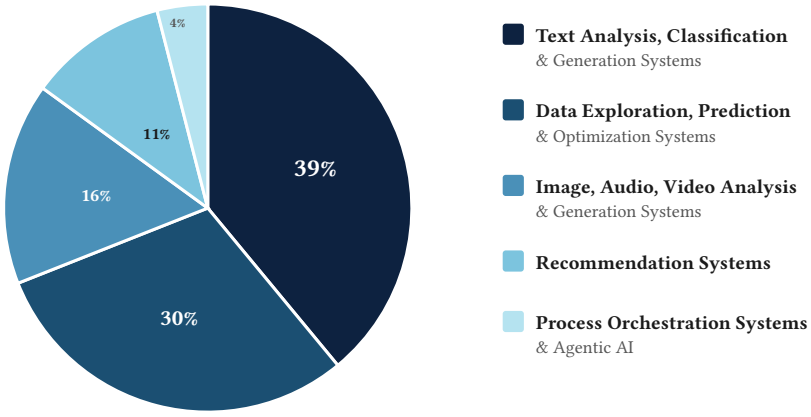

THE ADOPTION GAP

AI adoption in Italian companies (2025)

Source: AI Observatory, Politecnico di Milano — 698 companies surveyed

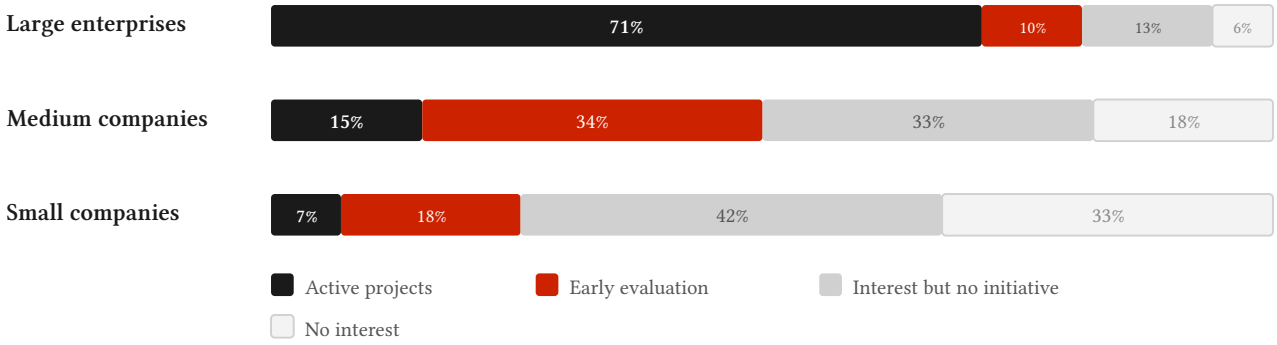

The gap is not about awareness — it's about access. Large enterprises have dedicated teams, cloud budgets, and consulting contracts. SMEs have spreadsheets, limited IT staff, and a justified fear of regulatory complexity.

The numbers speak clearly: **71%** of large enterprises have active AI projects — up from 59% in 2024. But only **7%** of small companies have active projects. **75% of small companies** haven't started anything yet. The gap is structural: large companies have internal teams, dedicated budgets, and the capacity to experiment. SMEs — which represent 95% of Italy's business fabric — remain largely on the sidelines.

> **WHAT THIS MEANS FOR YOUR BUSINESS**
>
> If you're an SME that hasn't started with AI, you're in the majority — but the window is closing. The businesses that move now, with compliant, private infrastructure, will set the standard. Those that wait will face both competitive disadvantage and regulatory pressure.

Germany, France, Austria and Switzerland

Eurostat's ICT Enterprise survey (isoc_eb_ai) provides the only directly comparable data across EU countries. The picture is clear — and Italy trails its peers.

GERMANY: LARGEST EU MARKET, MITTELSTAND UNDER PRESSURE

Bitkom estimates the German AI market at **over €9 billion in 2025**. Eurostat confirms 57% of large enterprises use AI, above EU average. But the Mittelstand struggles: only 23.1% of small companies have adopted. The governance gap is acute — only **23% of companies have formal rules for GenAI usage** (Bitkom).

FRANCE: MASSIVE POLITICAL BETS, LATE ADOPTION

France committed the most: **€109 billion in announced AI investments** at the Paris AI Summit, Bpifrance deploying €10B by 2029, Mistral AI as Europe's sovereign champion (€11.7B valuation, open-weight models). Yet the Cour des Comptes found that the **"expected massification has not taken place"**. Eurostat confirms: only 15% of small French enterprises use AI.

AUSTRIA & SWITZERLAND: THE SURPRISE LEADERS

Austria is the standout: **30% overall adoption, 68.3% of large enterprises** — highest among the four. Switzerland (not in Eurostat) reports **34% of SMEs consciously integrate AI** (KMU.admin.ch), up from 22% in 2024. Both punch above their weight, but gaps in strategy persist.

Country	Small (10–49)	Medium (50–249)	Large (250+)	All (10+)
Austria	**26.2%**	**44.6%**	**68.3%**	**30.0%**
Germany	23.1%	35.6%	57.0%	26.0%
EU 27	*17.0%*	*30.4%*	*55.0%*	*20.0%*
France	15.0%	30.8%	58.0%	18.2%
Italy	14.2%	27.6%	53.1%	16.4%

Source: Eurostat, isoc_eb_ai (2025) — % of enterprises using at least one AI technology. Switzerland not included (non-EU); Swiss data: KMU.admin.ch.

> *The pattern is identical across Europe: large enterprises adopt, SMEs hesitate, employees use ChatGPT anyway. The solution is the same everywhere — private, compliant AI.*

Shadow AI: the risk hiding in every European office

According to the AI Observatory at Politecnico di Milano, **47% of Italian workers** already use AI tools at work. But only 19% use exclusively company-provided tools. The rest use ChatGPT, Claude, or other tools with personal accounts — **Shadow AI**. And the pattern repeats in every country analysed: in Germany, only 23% of companies have formal GenAI rules. In Switzerland, only 24% offer mandatory AI training.

Shadow AI in European businesses

47% of workers use AI — but who controls it?

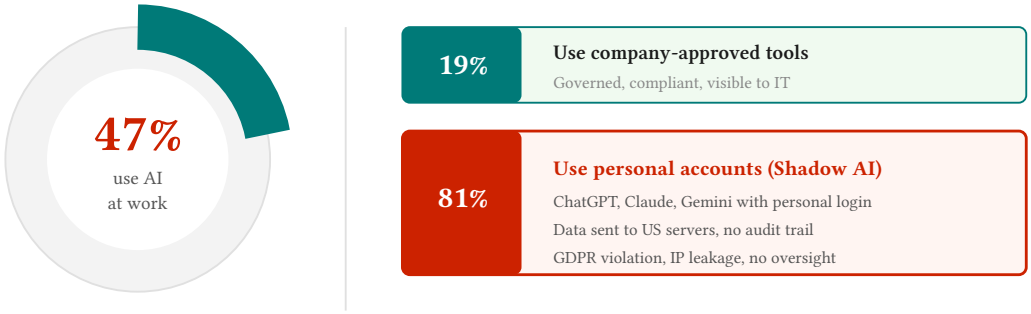

Source: AI Observatory, Politecnico di Milano (February 2026)

> *Every prompt sent to public ChatGPT passes through servers in the United States, where it can be used to train future models.*

For European businesses, this means simultaneous exposure on four fronts:

- **GDPR violation** — Personal data processed outside EU, without adequate safeguards
- **Intellectual property leakage** — Prompts may be used to train future models
- **AI Act non-compliance** — No audit trail, no risk assessment, no human oversight
- **Uncontrolled Shadow AI** — No visibility into what data employees are sharing

24% of large companies have already banned unauthorised GenAI tools. But banning is not enough — you need to **offer a private, controlled alternative**.

THE REAL COST OF SHADOW AI

Shadow AI isn't just a compliance risk — it's an operational blind spot. When employees use personal AI accounts, the company loses visibility into how decisions are being made, what data is being shared, and whether outputs are reliable. The solution isn't prohibition. It's giving employees a tool that's just as good — but private.

The untapped potential: the chart every CEO should see

Anthropic's research offers a global perspective on the potential still to capture. The chart below compares, for each occupational category, the **theoretical AI coverage** with **actual observed usage**:

Theoretical AI capability vs actual usage by sector

Source: Massenkoff & McCrory — Labor market impacts of AI, Anthropic (2026)

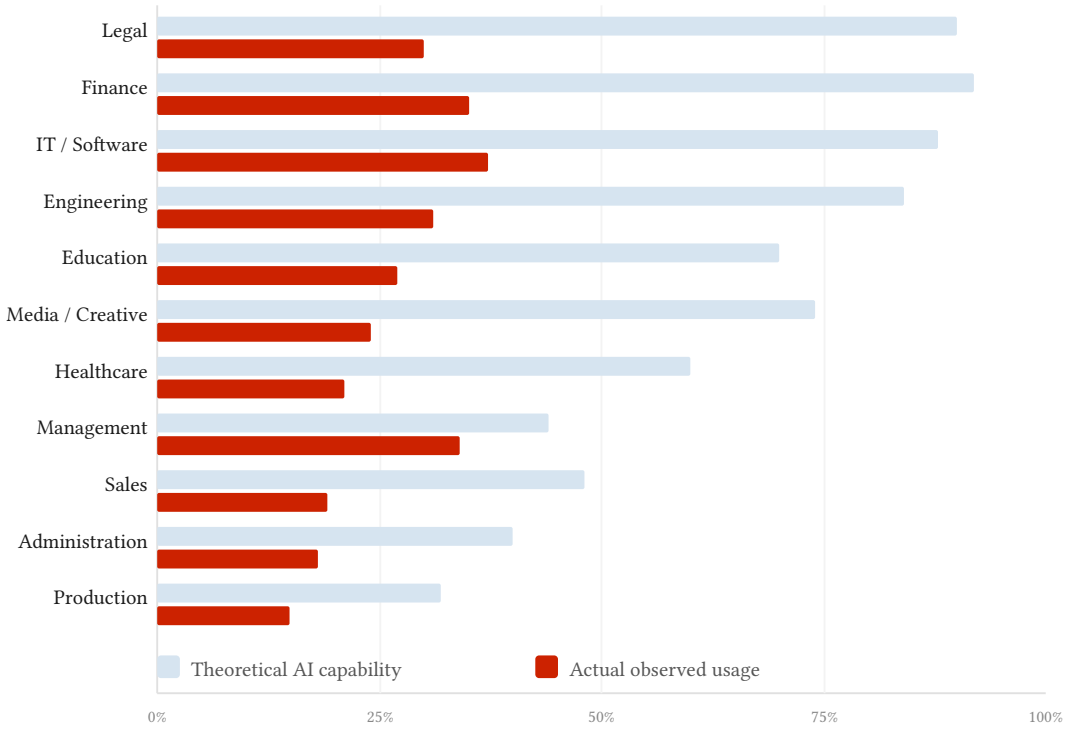

The gap is widest in **legal, finance, engineering, and IT** — precisely the sectors that handle the most sensitive data and therefore *cannot* use public cloud AI tools.

This creates a paradox: the businesses that would benefit most from AI are the ones least able to use it — because the available tools don't meet their privacy and compliance requirements.

> *The AI revolution isn't being held back by technology. It's being held back by trust — and by tools that don't deserve it.*

The solution isn't to wait for better regulation or cheaper cloud services. It's to deploy AI *on your terms*: on-premise, open-source, fully under your control.

THE OPPORTUNITY FOR EUROPEAN SMES

The gap between AI capability and actual adoption represents a massive competitive opportunity. SMEs that deploy private, compliant AI now will capture the productivity gains while larger competitors remain stuck in cloud vendor contracts and regulatory uncertainty.

CASE STUDY

From 50 days to 1.5: how we automated data migration with AI

How T&B Associati cut data entry from 50 man-days to 1.5 using HT-X's PRISMA platform. 52,000 hours processed, 180% ROI.

By Stefano Puissa, MBA — CEO, T&B Associati

At T&B Associati we specialise in complex project management. Our team coordinates resources across multiple assignments simultaneously, relying on project management software to plan activities, allocate resources, and track hours worked.

Like many European SMEs, we had a critical pain point: **data arrived in Excel format.**

THE PROBLEM

Our consultants, partners, and external collaborators reported their hours in Excel spreadsheets — each with their own format, structure, and naming conventions. Transferring this data into the centralised management system required:

- Manual interpretation of each spreadsheet
- Field-by-field data entry
- Cross-referencing with project codes
- Error correction and validation

The volume: **52,000 hours of work data** across dozens of projects. The time required: approximately **50 man-days** — nearly two and a half months of full-time work by a dedicated resource.

THE SOLUTION: PRISMA BY HT-X

We engaged HT-X, who deployed their PRISMA platform to automate the entire pipeline:

1. **Analysis phase** — The AI learned to recognise the structure of each Excel format, mapping columns to our management system's data model.
2. **Transformation phase** — Data was cleaned, normalised, and validated automatically, flagging anomalies for human review.
3. **Loading phase** — Clean data was imported directly into the management system, with full audit trail.

All processing happened **on-premise**. No client data left our infrastructure.

T&B Associati: before and after AI automation

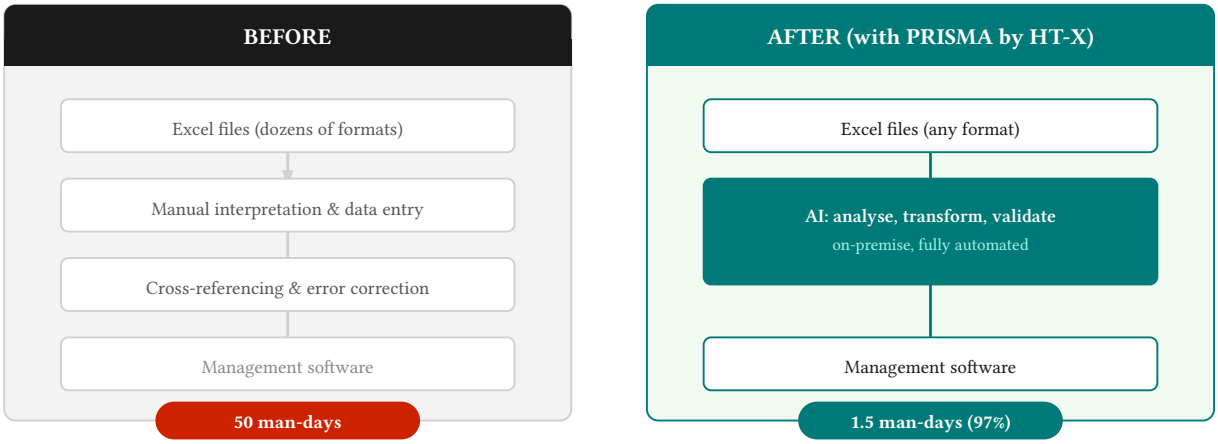

> *The AI didn't just copy data faster — it understood context, resolved ambiguities, and flagged what needed human attention.*

The investment paid off immediately, and it keeps paying off: the AI *learned* the patterns and now handles new spreadsheet formats with minimal reconfiguration. Every quarter, we save weeks of manual work.

RESEARCH

KOI: when AI helps doctors excel at anaesthesia risk classification

An on-premise AI system reduces ASA-PS classification errors by 89%. The story of a product born from research with Centro Ortopedico di Quadrante (Ramsay Santé).

By Francesco Menegoni, PhD

Every time a patient needs surgery, an anaesthetist assesses their health status and assigns a score: the **ASA-PS classification** (American Society of Anesthesiologists Physical Status). It's one of the most widely used systems in medicine — for over 80 years.

The problem? Doctors disagree. Studies on hundreds of anaesthetists show the correct classification is assigned only 70% of the time. In a third of assessments, consensus isn't even reached.

It's not a competence issue: it's a problem of *inherent variability* in a system based on subjective judgements.

THE RESEARCH: 11 MODELS, 20 CASES, 2 LANGUAGES

HT-X partnered with Centro Ortopedico di Quadrante, part of the Ramsay Santé group, to investigate whether large language models could bring consistency to this critical clinical task.

The study tested 11 different AI models on 20 real clinical cases, in both Italian and English. Each model received the same clinical information that a human anaesthetist would review.

ASA-PS classification accuracy: AI vs human doctors

20 real clinical cases, 11 AI models tested

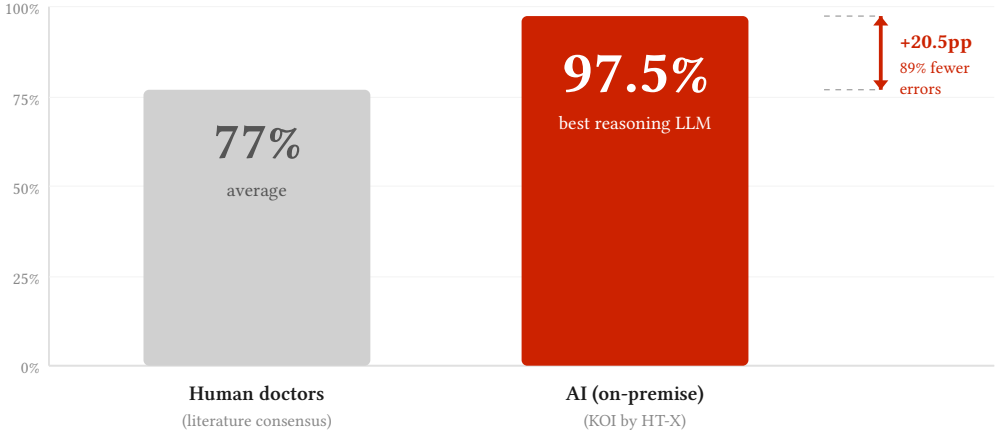

Study conducted with Centro Ortopedico di Quadrante (Ramsay Santé) — peer review submitted

The critical finding: **reasoning-capable LLMs** — models that show their step-by-step thinking — achieved the highest classification accuracy, outperforming traditional LLMs and providing doctors with a level of consistency that is difficult to maintain across hundreds of cases. The AI doesn't "guess" based on pattern matching; it reasons through clinical evidence systematically.

WHY ON-PREMISE, NOT CHATGPT

The clinical data involved in ASA-PS classification is among the most sensitive that exists: complete patient health records. Sending this to cloud-based AI services would violate both GDPR and basic medical ethics.

HT-X's testing showed that **on-premise open-source models achieved identical performance to cloud-based proprietary models**. There was no accuracy trade-off for choosing privacy.

> *The AI doesn't replace the anaesthetist — it gives them a second opinion that's consistent, evidence-based, and never tired.*

FROM PAPER TO PRODUCT

The research has been submitted for peer review to *Informatics in Medicine Unlocked*. Simultaneously, HT-X developed **KOI** — a productised version that hospitals can deploy on their own infrastructure, integrated into existing clinical workflows.

KOI represents a model for how AI should enter healthcare: born from rigorous research, deployed privately, and designed to *support* — not replace — clinical judgement.

FOR HEALTHCARE ORGANISATIONS

If your facility performs surgical procedures and wants to explore AI-assisted clinical assessment, KOI can be deployed as a pilot within your existing infrastructure. No patient data leaves your systems. Contact us at ht-x.com.

TECHNOLOGY

Fine-tuning: how to teach AI your business language

A guide to LLM fine-tuning. What it is, when you need it, and why it's the step that turns a generic model into a real working tool.

By Ivan Buttignon, BSc

ChatGPT, Claude, and Gemini are powerful models, but generic. They know everything about everything — and nothing about your company. They don't know your terminology, procedures, communication tone, or document structure.

The result? Approximate answers requiring constant corrections. Increasingly long prompts to explain context. Inconsistent results from one day to the next.

Fine-tuning solves this at the root: instead of explaining what to do every time, you teach the model how — once and for all.

The three levels of AI customisation

Start simple, add complexity only when needed

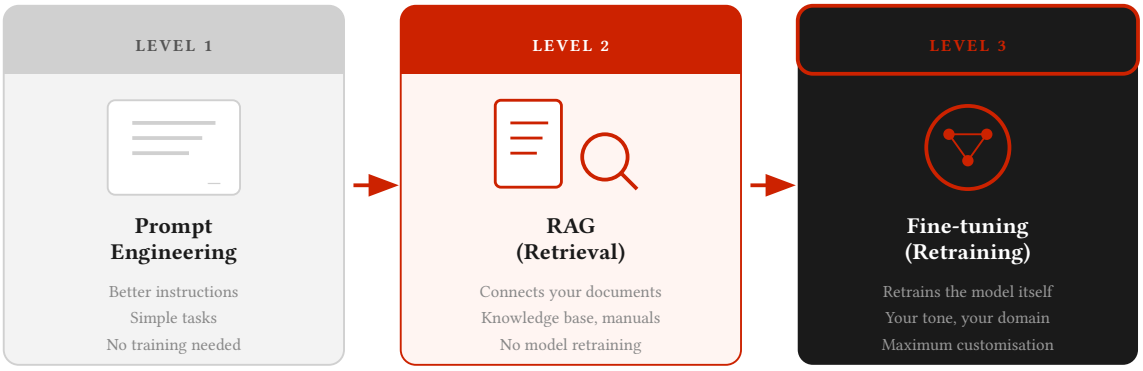

Start at Level 1. Move to the next only when you've hit the limits of the current one.

Most businesses should start with prompt engineering, add RAG when they need access to internal documents, and consider fine-tuning when they need the model to *behave* differently — not just *know* more.

THE TECHNIQUES

Not all fine-tuning is equal. Three approaches exist, with very different resource requirements:

- **Full Fine-Tuning** — Retrains every parameter. Maximum quality, but requires significant GPU resources and large datasets. Suited for fundamental domain shifts.

- **LoRA** (Low-Rank Adaptation) — Trains only small adapter layers. 90% of the quality at 10% of the cost. The sweet spot for most business applications.
- **QLoRA** — LoRA on quantised models. Runs on consumer-grade GPUs. Ideal for prototyping and smaller deployments.

WHAT YOU GET IN PRACTICE

A fine-tuned model can:

- Answer customer queries in your exact tone and with your product knowledge
- Analyse documents using your classification criteria
- Generate reports in your house style
- Route support tickets using your category system
- Write technical documentation matching your engineering standards

> *Fine-tuning doesn't make the AI smarter — it makes it yours.*

WHY FINE-TUNING MUST HAPPEN ON-PREMISE

Fine-tuning requires feeding the model your most valuable data: internal documents, customer interactions, proprietary processes. This is precisely the data that must *never* leave your infrastructure.

With ORCA by HT-X, fine-tuning happens entirely on your servers. The training data, the resulting model, and all intermediate artefacts remain under your control. No cloud provider sees your data. No third party benefits from your investment.

GETTING STARTED

HT-X offers a structured six-step fine-tuning journey: assessment, data preparation, training, validation, deployment, and monitoring. The entire process typically takes 4–8 weeks, depending on data volume and complexity.

REGULATION

AI content labelling: what changes on 2 August 2026

The European Commission's Code of Practice on AI content marking. Watermarks, metadata, and icons become mandatory.

By Francesca Spazzali, MSc

AI content labelling: key dates

Watermarks, metadata (C2PA) and visible labels become mandatory for all AI-generated public content

On 5 March 2026, the European Commission published the second draft Code of Practice on marking and labelling AI-generated content. This isn't an academic draft: it's the document that will define **how businesses must mark text, images, audio, and video produced with generative AI**.

The final version arrives in June 2026. The obligation kicks in on **2 August 2026**.

WHAT THE AI ACT SAYS

Article 50 of the AI Act establishes that AI-generated content must be detectable as such. The Code of Practice translates this principle into technical requirements:

1. **Machine-readable marking** — Secured metadata embedded in the file (C2PA standard for images/video, IPTC for text)
2. **Watermarking** — Invisible markers that survive common transformations (cropping, compression, format conversion)
3. **Visible labelling** — Icons or text labels for content presented to end users

The two-layer approach — metadata + watermark — ensures marking survives even if a file is shared, edited, or re-uploaded.

WHAT THIS MEANS FOR BUSINESSES

Any company using AI to generate content for external communication must prepare:

Content type	Obligation
Marketing images (AI-generated)	Metadata + watermark + visible label
AI-written text (blogs, reports)	Metadata marking (IPTC)
AI-generated audio/video	Watermark + metadata
Internal documents	No obligation (not public-facing)
AI-assisted (human-edited)	Depends on degree of AI contribution

THE PRIVACY ANGLE

Here's the catch: if your business generates AI content through cloud services (ChatGPT, Midjourney, DALL-E), the *provider* is responsible for embedding the marking. But your content — including drafts, prompts, and context — is processed on their servers.

With on-premise AI, you control the entire pipeline: generation, marking, and distribution. The content never leaves your infrastructure, yet you can still apply compliant metadata and watermarks.

HOW TO PREPARE

Five steps before August 2026:

1. **Audit** — Inventory all AI-generated content in your organisation
2. **Classify** — Determine which content is public-facing and needs marking
3. **Tooling** — Ensure your AI tools support C2PA metadata and watermarking
4. **Workflow** — Integrate marking into your content production pipeline
5. **Training** — Ensure teams understand when and how marking applies

> **TIMELINE**
>
> June 2026: Final Code of Practice published. August 2026: Obligation enters into force. Businesses using AI for content creation should begin preparation now.

GUIDE

Private ChatGPT: the complete guide for European businesses

How to use ChatGPT-level AI in your business without risking your data. A practical overview for decision-makers.

By Fabio Tiberio, BSc

According to the AI Observatory of Politecnico di Milano, 47% of Italian workers already use AI tools at work. But only 19% use exclusively company-provided tools. The rest use ChatGPT, Claude, or other tools with personal accounts — **Shadow AI**.

Every prompt sent to public ChatGPT passes through OpenAI's servers in the United States. For European businesses, this creates four simultaneous risks:

- **GDPR violation** — Personal data processed outside EU, without adequate safeguards
- **Intellectual property leakage** — Prompts may be used to train future models
- **AI Act non-compliance** — No audit trail, no risk assessment, no human oversight
- **Uncontrolled Shadow AI** — No visibility into what data employees are sharing

THE SOLUTION: ON-PREMISE PRIVATE AI

A private ChatGPT alternative runs the same calibre of AI models — but on your infrastructure:

Feature	Public ChatGPT	ORCA (on-premise)
Data location	US servers	**Your servers**
GDPR compliance	Problematic	**By design**
AI Act readiness	Unclear	**Built-in audit trail**
Model choice	GPT only	**Llama, Mistral, DeepSeek…**
Fine-tuning	Limited	**Full control**
Cost model	Per-token	**Fixed infrastructure**
Vendor lock-in	High	**None (open-source)**
Offline capable	No	**Yes**

WHO NEEDS PRIVATE AI

Any business that handles:

- **Patient data** — Healthcare, clinical research, pharma
- **Financial data** — Banking, insurance, accounting
- **Client confidential information** — Legal, consulting, M&A
- **Intellectual property** — R&D, engineering, manufacturing
- **Customer personal data** — Any B2C business in the EU

THE OPEN-SOURCE ADVANTAGE

In 2026, open-source models have closed the gap with proprietary ones.

Llama 3, Mistral, DeepSeek R1, Qwen 3.5, and GLM 5 deliver performance comparable to GPT-4 and Claude — with full transparency, no usage fees, and complete data sovereignty.

The choice is no longer between "powerful but risky" and "safe but limited." Private, on-premise AI now delivers *both*.

> *Compliance shouldn't be the company's problem. It should be a feature of the solution.*

PRODUCT

ORCA

Your AI. Your data. Your infrastructure.

ORCA is HT-X's on-premise AI platform. It gives businesses access to the full power of large language models — without sending a single byte of data outside their infrastructure.

ORCA: on-premise AI architecture

Your data never leaves your infrastructure

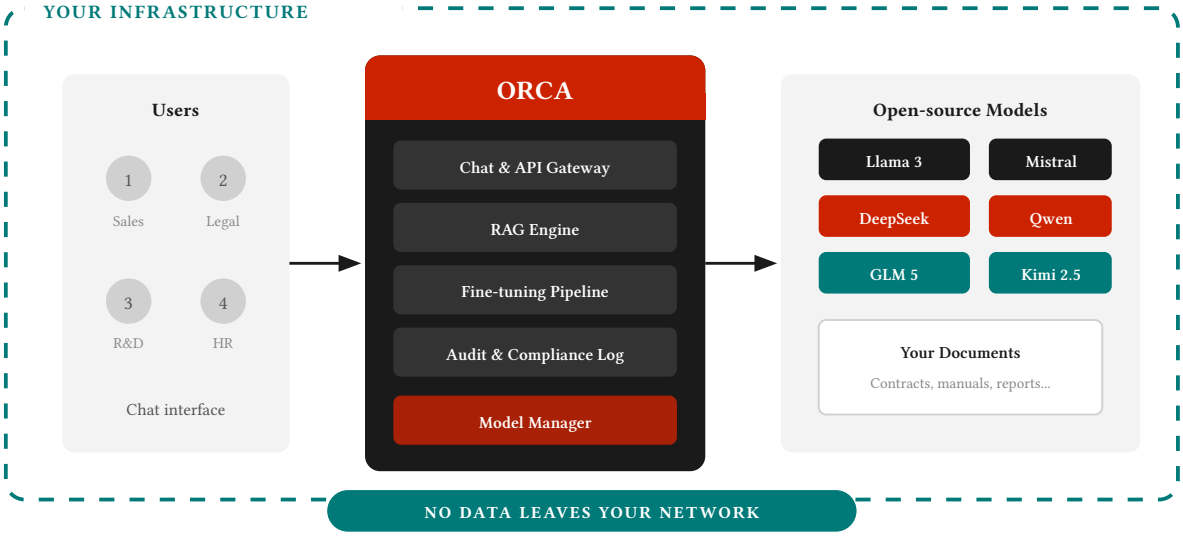

WHAT YOU GET

ChatGPT-like interface

Familiar chat experience for every employee. No training needed.

Multiple AI models

Run different models for different tasks. Switch freely, no vendor lock-in.

Document intelligence

Upload contracts, manuals, reports. Ask questions. Get answers with citations.

Full audit trail

Every interaction logged. AI Act compliant by design.

Fine-tuning ready

Train models on your data. Your terminology, your tone, your standards.

Offline capable

Works without internet. Ideal for air-gapped environments and classified data.

Ready to take control of your AI?

Visit **ht-x.com/orca** or find us at SMAU Paris — April 2026

Private AI Europe

Your AI. Your data.
Your rules.

---

privatechatgpt.eu

ht-x.com/orca

HT-X S.r.l. — Human Technology eXcellence
VAT IT01407090321 — Share capital €72,000

MEET US AT

SMAU Paris — April 15, 2026

Station F — 5, parvis Alan-Turing, Paris

AI is transforming business. Where do you stand?

Take the free assessment