

Industrial AI Audit Guide

The Industrial AI Audit Guide

How to assess your AI maturity, map your information system, identify your value pools and build an operational roadmap.

Taranis AI Guide — 2026 edition. Method aligned with the Industrial AI Audit page and the NIST AI RMF and IIA reference frameworks.

Why Audit Before You Invest

Most AI projects fail before they ever reach production. The figures are documented: **at least 30% of GenAI projects were abandoned after proof of concept by the end of 2025**, according to Gartner, "due to poor data quality, inadequate risk controls, escalating costs or unclear business value" [Gartner press release, 29/07/2024 (<https://www.gartner.com/en/newsroom/press-releases/2024-07-29-gartner-predicts-30-percent-of-generative-ai-projects-will-be-abandoned-after-proof-of-concept-by-end-of-2025>)]. And the trap does not stop there: **Gartner predicts that through 2026, organisations will abandon 60% of AI projects not supported by AI-ready data** — data aligned to a use case, governed at the asset level, supported by automated pipelines with quality gates [Gartner press release, 26/02/2025 (<https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk>)].

In other words: **a better model does not rescue bad data**. The audit exists precisely to avoid launching those projects, or to launch them on the right scope. It is the first step of any successful AI transformation — and it is what this guide helps you structure, whether you run the audit internally or bring in a third party.

Part 1 — The Diagnostic Framework

An industrial AI audit rests on three building blocks: assessing your **AI maturity**, mapping your **information system**, and identifying your **value pools**. This is the diagnostic phase — the first step of the Taranis method (diagnosis → reorganisation → automation).

1.1 Assess Your AI Maturity

Maturity is not measured by the presence of a chatbot or a POC. It is measured across four dimensions, aligned with the reference **NIST AI Risk Management Framework (AI RMF)**, published on 26 January 2023, which organises the audit around four functions: **Govern, Map, Measure, Manage** [NIST, AI Risk Management Framework (<https://www.nist.gov/itl/ai-risk-management-framework>)].

Dimension	Question to ask
Processes	Which business processes are candidates for AI? Are they documented, repetitive, high-variability?
Data	What data do you already produce (sensors, ERP, MES, CRM)? Is it accessible, clean, governed?
Governance	Who is responsible for what? Are the roles documented? Does a usage charter exist?
Skills	What AI skills do you have in-house?

The maturity signal is clear among large companies: **60% of organizations have deployed cross-functional AI governance to industrialize AI**, according to the second edition of KPMG's *Trends of AI 2026* study, based on 356 French decision-makers [KPMG, Trends of AI 2026 (<https://kpmg.com/fr/fr/media/press-releases/2026/01/adoption-ia-entreprises-francaises.html>)]. Among SMEs and mid-caps, Bpifrance Le Lab's survey finds that **26% of companies with 10–5,000 employees already use generative AI**, while 28% more plan to [Bpifrance Le Lab, via independant.io (<https://independant.io/statistiques-ia-intelligence-artificielle>)]. The skills gap is the first lever: the French Senate's report *Enterprise 5.0* (April 2026) notes that **76% of French people have received no AI training** [French Senate, report n°572, 04/28/2026 (https://www.senat.fr/rap/r25-572/r25-572_mono.html)].

1.2 Map Your Information System

This is the step most companies dread because it reveals "shadow AI": tools adopted by teams without central validation. The **Map** function of the NIST AI RMF consists of "identifying risks and contributing factors" [NIST AI RMF (<https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>)].

The inventory, system by system:

- **Purpose:** what the system does, which decision it influences.
- **Data:** what data it uses, where it comes from, whether it is representative.
- **Users:** who uses it, and who is affected by its outputs.
- **Lifecycle:** who deployed it, when, with which version.

Coming out of this step, you have the answer to "how many AI systems do you use?" — often more than you think. Without an inventory, no risk classification; without classification, no compliance plan.

1.3 Identify the Value Pools

Not all processes are equal. Value pools concentrate on **repetitive, manual, high-variability processes** — where AI delivers measurable value, not where it "looks good on a PowerPoint slide." Four domains concentrate most industrial deployments, with documented returns:

Domain	Documented return	Source
Predictive maintenance	–30 to 45% unplanned stoppages	McKinsey Global Institute, IoT 2015
Predictive maintenance	–25 to 30% maintenance costs	Field data, Deloitte Smart Factory studies
Vision-based quality control	Inspection accuracy from ~80% to 99.86%	Sundaram et al., 2023
Production output	+10 to 20%	Deloitte 2026 Smart Manufacturing Study

The rule of thumb: **start with the processes where the data already exists and the return will be fastest.** An AI project that fails rarely does so for technical reasons — misalignment with business needs, inaccessible data and absent ROI measurement are the dominant causes.

Part 2 — Prioritise the Use Cases

Once the diagnosis is done, you must decide. The **Manage** function of the NIST AI RMF consists of allocating risk-management resources to mapped and measured risks [NIST AI RMF (<https://airc.nist.gov/airmf-resources/airmf/5-sec-core/>)]. In practice, prioritisation runs on two axes:

- **Value**: estimated time savings, error reduction, projected ROI.
- **Effort**: technical feasibility, quality of available data, budget, required team.

The result is a selection of the **5 to 10 use cases with the best value/effort ratio** in your specific context. This is what turns the investment decision from intuition into a documented choice: you know exactly what you stand to gain before investing a single euro.

Each use case is then costed: time savings, error reduction, required budget, required team. You get a clear view of which projects to launch first, which to deploy next and which to drop.

Part 3 — The Roadmap

The output of the audit is a **costed roadmap**: which systems, which gains, which risks, which costs, in which order. A realistic timeline over **6 to 18 months**, split into three phases consistent with the Taranis method (diagnosis → reorganisation → automation).

Phase 1 — Diagnose (months 1-2)

Before any technology purchase, take the stock described in Part 1: what data you already produce, which processes are candidates, what skills you have in-house. Bpifrance offers a structuring program: the **AI & Industry Accelerator** (18 months, total cost €71,800 ex. tax, with €25,400 covered by Bpifrance via the *Osez l'IA* plan) including 12 days of Data AI Industry diagnosis [Bpifrance, AI & Industry Accelerator (<https://www.bpifrance.fr/catalogue-offres/accélérateur-ia-industrie>)].

Phase 2 — Start with a quick-ROI use case (months 3–6)

Choose a first project with high return and low risk: predictive maintenance on one critical machine, or automated quality control on one line. First results should be visible in 3 to 6 months to build confidence and justify the next steps.

Phase 3 — Industrialize and govern (months 6–18)

Once the first case is validated:

- Set up **AI governance** (usage charter, system register, responsibilities).
- Deploy **cross-functional governance** (KPMG confirms 60% of organizations already do).
- Train your teams: the Senate report notes that **76% of French people have received no AI training** — this is the primary competitiveness lever for companies investing in upskilling.

Anticipate Compliance

The audit is also the natural entry point to compliance. The EU AI Act (Regulation (EU) 2024/1689) is applying progressively: the transparency obligations (art. 50) have applied since **2 August 2026**; the obligations for autonomous high-risk systems arrive on **2 December 2027**, and those for high-risk systems embedded in regulated products on **2 August 2028** — a timetable postponed by the Digital Omnibus, in force since 27 July 2026 [EUR-Lex (<https://eur-lex.europa.eu/eli/reg/2026/1744/oj>)]. Fines can reach **7% of worldwide annual turnover** for prohibited practices [Article 99, AI Act (<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689>)]. A documented audit is what lets you answer the first question any auditor or authority will ask: **do you know what you have?**

Do It Yourself or Hire a Third Party?

The method above can be run internally. The reference frameworks are public (NIST AI RMF, IIA). It is the logical first step for an SME or mid-cap company: you learn where your data and risks are before you invest.

A provider becomes relevant in three situations:

1. **Independence:** auditing a system your own team built carries a judgement bias. An external eye — NIST calls it "effective challenge," the critical questioning of design decisions by experts with the authority to make changes [NIST AI RMF Playbook (<https://airc.nist.gov/airmf-resources/playbook/govern/>)] — uncovers what internal complacency masks.
2. **Regulatory framework:** for AI Act high-risk systems, documentation and conformity assessment demand rigorous methods most internal teams do not yet have in place.
3. **Speed:** a costed roadmap in 6 weeks, with prioritised use cases and projected ROI, requires a proven method and comparable deployments.

Rule of thumb: start internally to learn, externalise for independence, compliance and speed.

Conclusion: an Audit Is Judged on the Decisions It Triggers

An AI audit is not another report. It is the instrument that turns your AI systems from a collection of scattered initiatives into a **managed portfolio**: named responsibilities, known data, identified value pools, documented investment decisions and AI Act compliance in preparation.

The cost of an audit without a method is funding the 30% of projects abandoned after proof of concept and the 60% killed by data that was never ready — in each case, money, time and team trust. The benefit of an audit with a method is launching only what pays, and being able to prove it.

Start with the first step, the one that costs nothing: **map your systems**. How many AIs do you actually use, for which decisions, with which data?

👉 **Request an AI audit** — We map your processes, quantify the gains, prioritise the use cases and hand you a costed roadmap in 6 weeks.

Guide published in September 2026 by Taranis AI. Method aligned with the NIST AI Risk Management Framework (AI RMF 1.0, January 2023) and the IIA audit framework; project-failure figures from Gartner (press

releases of 29/07/2024 and 26/02/2025); adoption data from KPMG (Trends of AI 2026), Bpifrance Le Lab, French Senate (Enterprise 5.0, report n°572); industrial returns from McKinsey Global Institute (IoT 2015), Deloitte Smart Manufacturing and Sundaram et al. (2023); AI Act deadlines verified on EUR-Lex (Regulation (EU) 2024/1689, Digital Omnibus (EU) 2026/1744).