TERRIA CONSEIL

2026 MARCH SYNTHESIS

WHITE PAPER

AI in business, understanding the gap.

Study conducted with 135 decision-makers and IT professionals in France and Belgium

Co-written by Terria Conseil founders

EXECUTIVE SUMMARY

Our synthesis in one page

Generative AI has spread rapidly through businesses. Use cases have frequently preceded governance frameworks, driven by individual initiatives. Our initial focus was solely on security issues surrounding this topic (e.g., usage oversight, data management). Our partners confirmed these concerns and also gave us a broader view of the challenges posed by AI.

The purpose of this summary is to present, in a framework useful to those interested in the subject, what we gathered from our partners (i.e., 135 decision-makers and IT professionals in France and Belgium). We cross-referenced this information with academic and institutional studies to support the analysis. This document is a consolidation designed to provide an overview of the current situation and help you find your footing within this AI-driven dynamic. We have no ambition to revolutionize your understanding of AI challenges, provide an exhaustive view nor to offer a magic formula. Our aim is to provide this analysis grounded in real-world feedback, free from both alarmism and optimism.

Key take-aways :

1. AI entered companies through curious employees, outside of formal projects. The most common uses remain fairly superficial (e.g., advanced office productivity, an enhanced search engine), and their adoption is uneven across the workforce.
2. The pressure to "not miss the AI wave" is pervasive and constant. It comes from the market, vendors, consultants, and the media, and crystallizes into urgent requests from business units. This dynamic accelerates decisions and undermines sustainable adoption.
3. Executive awareness is real but partial. Several leaders treat AI as just another technology cycle, as they once did with cloud or SaaS, assuming the same playbook will apply. They do not frame AI as a technological disruption that needs to be fully integrated into their company's strategy.
4. IT teams are facing mounting pressure. Vibe coding produces micro-applications that are hard to identify and maintain. Agentic AI is appearing on the horizon with risks that organizations still struggle to grasp.
5. Support functions are seeing their role shift within the company. AI tools give employees plausible and immediate answers on expert topics that were historically the domain of these functions.
6. Cognitive effects are beginning to be documented, such as the illusion of competence or a reversal of the Dunning-Kruger effect in the presence of AI. The more we use it, the less accurately we assess our own level of competence and the quality of what it produces.
7. The pace of the AI ecosystem is creating tension with the pace of organizations. A new innovation is announced every week, while a company needs several months for a successful transformation.

The main avenues for action :

Three options emerge from our discussions:

- Situate before acting: map actual usage, including Shadow AI and Grey AI.
- Rethink governance so that it keeps pace, applying principles of subsidiarity.
- Train people not only to use AI tools, but also to evaluate the quality of the work they produce.

ABOUT TERRIA CONSEIL :

Terria Conseil is a company founded in 2026 by three AI optimists. We believe every business should be able to navigate the AI transition on its own terms. Our platform is built as an alternative to traditional consulting, empowering companies to internalize this critical topic, equip their teams and approach this shift with confidence.

CONTENTS

A three-part synthesis

YOU ONLY HAVE 10 MIN ?

Focus on the sections marked with this symbol LES ESSENTIELS

00 - CONTEXT

Why this study, methodology & hypotheses

WHY THIS STUDY

We sometimes have the impression that artificial intelligence emerged spontaneously. When we step back from the daily announcements, we realize that these technologies have existed for several decades. They were present in laboratories and large corporations, but barely visible in our everyday lives. Indeed, these tools were complex and costly to deploy in terms of skills, data, and infrastructure. They were therefore limited to use cases with very high added value.

The launch of ChatGPT 3.5 in November 2022 marked a turning point in the dynamics of AI technologies. They became easily accessible at low cost and without prior expertise, at least in appearance. OpenAI's success clearly illustrates this paradigm shift, with 100 million active users in two months. By comparison, it took Facebook four years to reach that threshold, and TikTok two years. From that moment on, a groundswell was set in motion. A shift that has shaken many pillars of our lives and created new personal and professional habits.

In 2026, as we launch our entrepreneurial venture, we want to do our part in this ongoing paradigm shift. We aim to make a contribution that reflects who we are, something useful and pragmatic, without over-promising. We are convinced that artificial intelligence is a double-edged sword for businesses: both an opportunity not to be missed and a catalyst for risk.

We launched this study to understand how companies are navigating this turning point. Given our respective backgrounds, we came in with our own biases and hypotheses. These shaped our approach, but through our conversations with the people we interviewed, we broadened our analysis in order to present the lessons we drew from it as faithfully and meaningfully as possible.

A CO-FOUNDER'S PERSPECTIVE

Despite my passion for technical innovation, I am by nature skeptical when it comes to big technological announcements. I followed blockchain, VR, the metaverse... more out of curiosity than out of any conviction that I was witnessing a major technological breakthrough.

But I will long remember the evening in December 2022 when I opened ChatGPT for the first time. I raised an eyebrow and felt an excitement I hadn't experienced in a long time. A sleepless night followed, along with the conviction that I needed to follow this technology closely to see where it would take us.

That conviction is what gave rise to Terria Conseil and this study

00 - CONTEXT

Why this study, methodology & hypotheses

METHODOLOGY

This study is based on two complementary approaches, conducted with **135 decision-makers and IT professionals in France and Belgium.**

40

INTERVIEWS

Live interviews with
CIO/IT Director, CISO, IT decision-makers

95

SURVEYS

Responses to our survey with closed question.

The questionnaire provided quantitative data on the initial angle of our study. The open-ended interviews allowed us to gather the details of real-life situations and to go beyond our initial hypotheses.

Here are some additional statistics about the panel:

Respondents function

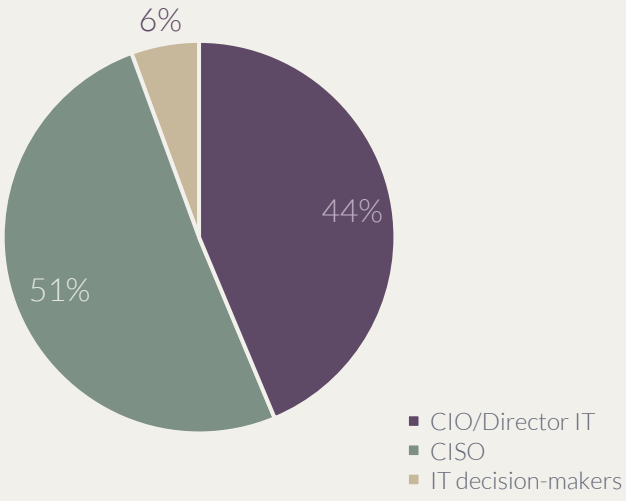

Companies size

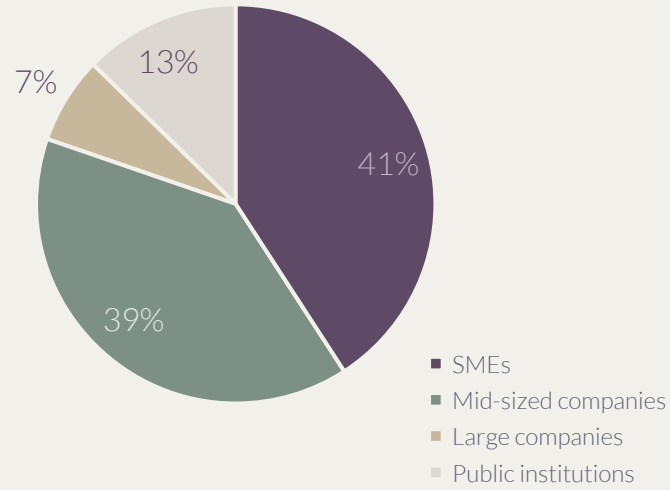

As our study broadened over time, we enriched our analysis by reading and integrating academic and private research studies to deepen the insights gathered during the interviews. For each section, we have compiled the studies we found most relevant and on which we drew.

INITIAL HYPOTHESES, EXPANSION & LIMITS

WHERE DID WE START?

We have a strong interest in cybersecurity and digital sovereignty. Our starting angle was heavily shaped by these domains. In our experience, securing a company's Information System was already a major challenge before AI arrived– Shadow IT, technical debt, complex operations and security maintenance, data leaks, etc. The explosion of AI use has created new attack vectors that increase risks and further complicate IT security. This trend is confirmed by studies from ANSSI and CESIN.

We began our inquiry aiming to answer three questions:

1. How are companies governing these new uses?
2. How are they managing data-related risks (e.g., detection, blocking, anonymization)?
3. Do they have the tools to address them?

Our initial hypothesis was that AI is a major security challenge for businesses.

THE EXPANSION WE MADE

As we entered the interview phase, we quickly realized that, even if our initial hypothesis was valid, it was above all reductive of the real challenges businesses face with the arrival of AI. Ultimately, the question is not solely about security.

Issues related to Shadow AI, that is, the use of AI tools not validated by the company, data leaks, and new forms of attacks are very real. However, **the central issue is one of absorption:** organizations are structurally not ready to integrate a technology that evolves faster than they can adapt.

This gap between the pace of AI and businesses' capacity to absorb it is ultimately at the heart of the challenge companies face.

LIMITATIONS OF OUR STUDY

Before you continue reading, it is important for us to clarify the limitations we have identified for this study:

- The broadening of scope could only be applied to the interviews, which means **the survey is only usable for the security-related portion of the study.**
- The insights we drew are based on feedback from 40 people. **This study therefore represents a sample** of what companies are experiencing with the arrival of AI.
- This analysis **does not aim to provide absolute truth** or an omniscient view of the business landscape.
- We supplemented our analysis with **studies from public or private institutions** that may **have their own interests, which introduces potential biases.**
- Our panel introduces a bias linked to a **predominantly IT perspective** on these AI challenges.

01 — AI IMPACT ON CYBERSECURITY

Our initial hypothesis & the reality of our conversations

1.1 — SECURITY ALREADY UNDER STRAIN BEFORE AI

Our initial hypothesis was rooted in an observation made over several years in businesses. Even before generative AI arrived, organizations were already struggling to master their digital perimeter.

This observation is not just our personal feeling or that of interviewees: reference studies confirm it.

The ANSSI Cyberthreat Landscape 2025 is a good example: the level of cyber threat remains high in France and spares no one– individuals, micro-businesses, SMEs, mid-sized companies, etc. The number of incidents linked to data exfiltration has significantly increased compared to 2024 (+50% of incidents reported in one year). The boundaries between cybercriminals and state actors continue to erode. Attackers are specializing, sharing tools, and exploiting vulnerabilities in insufficiently supervised and prepared environments.

+200

data leaks referenced in France over the first four months of 2026, **i.e. more than one leak per day**.

Source : bonjourlafuite.eu.org (recensement communautaire, indicateur de tendance)

The companies we interviewed confirmed these difficulties: Shadow IT, accumulated technical debt, complex or deprioritized operations and security maintenance, insufficient access segmentation, etc.

1.2 — GENERATIVE AI: AN AGGRAVATING FACTOR

The arrival of generative AI did not create a new problem. It intensified an existing situation across three simultaneous and interdependent axes:

01

Accessibility

AI tools are accessible, often free, and easy to use.

02

Promises

The AI ecosystem continuously makes announcements that could revolutionize how we work.

03

Risks

Existing risks are intensified and new ones are emerging.

1.2 — GENERATIVE AI: AN AGGRAVATING FACTOR

Accessibility within everyone's reach

Before generative AI, deploying a tool that transformed how people worked required a project, budget approval, and IT team involvement. AI tools have changed this logic: ChatGPT, Copilot, Gemini, and Claude have free or low-cost versions that change our work habits without any particular installation or need for validation. The promises and possibilities of these tools push employees to adopt them without waiting for the company or the IT department.

The balance of power is structurally *un*balanced. A CIO interviewed in our study described it this way:

> *"There are 30 of us in IT, and 3,500 of them [employees]. It's a race we've already lost."*

Some companies try to block these tools, but without success. Another account illustrates the "creativity" of users in the face of blocking:

> *"We tried banning ChatGPT. I caught an employee sending documents to their personal phone, using them in their personal ChatGPT, and sending the responses back to themselves via WhatsApp. We had a sense of security, and in the end, it was the exact opposite."*

Promises that make stepping back difficult

A large majority of the people we spoke with share the same observation: announcements and promises from AI vendors, combined with business units' fear of missing the turn, create a frantic race within companies. The potential gains from AI tools seem considerable and the demonstrations appear convincing. For business teams that may be under pressure or facing recurring problems, it can be hard not to be tempted by a sales pitch that also makes the real complexity of such a project disappear– IT integration, data management, maintenance, training, etc.

A CIO illustrated this for us:

> *"Every time I accompany a business team to a vendor meeting, I'm the one who throws a wrench in the works. I see the complexity behind the promises that seem unrealistic, so I break the myth of magic AI."*

A risk exposure that remains barely perceptible

We identified one final common theme in the feedback from interviewees: the explosion of AI-related risks. All our interviewees perceive it clearly, but it remains difficult to quantify, even for those whose job is to do exactly that. Some risks are directly tied to these new tools (e.g., prompt injection, model hijacking, deepfakes). Others emerge from their use (e.g., unintentional data exfiltration, loss of control over developed code, destructive actions by autonomous agents) and they amplify existing risks. The apparent security of the tools, the lack of collective hindsight on their impacts and the novelty of uses, too recent to be documented and objectively assessed, make risks hard to grasp for teams. This makes it difficult within organizations for those who want to mitigate to be heard.

1.3 — DATA FROM OUR SURVEY

The data from our survey confirms and refines this picture:

76%

have a policy regulating or restricting AI

And only 2% completely ban AI use

86%

believe that Shadow AI exists within their company

Of which 44% believe it is routinely done

58%

have no technical measures against data leaks

Raising awareness remains the primary lever for action

The key elements that stood out to us:

- The vast majority of the companies surveyed have an official position on AI (89%).
- The dominant posture is regulation through usage rules or limitation to approved tools (76%), some respondents have implemented a policy without restrictions (11%), and only 2% ban it entirely.
- Companies that ban AI have done so for ethical or social reasons.
- Shadow AI is a reality shared by almost everyone (86%).
- Only 5% of respondents admit having no visibility into their actual AI traffic.

The risk associated with data leaks is the number one risk.

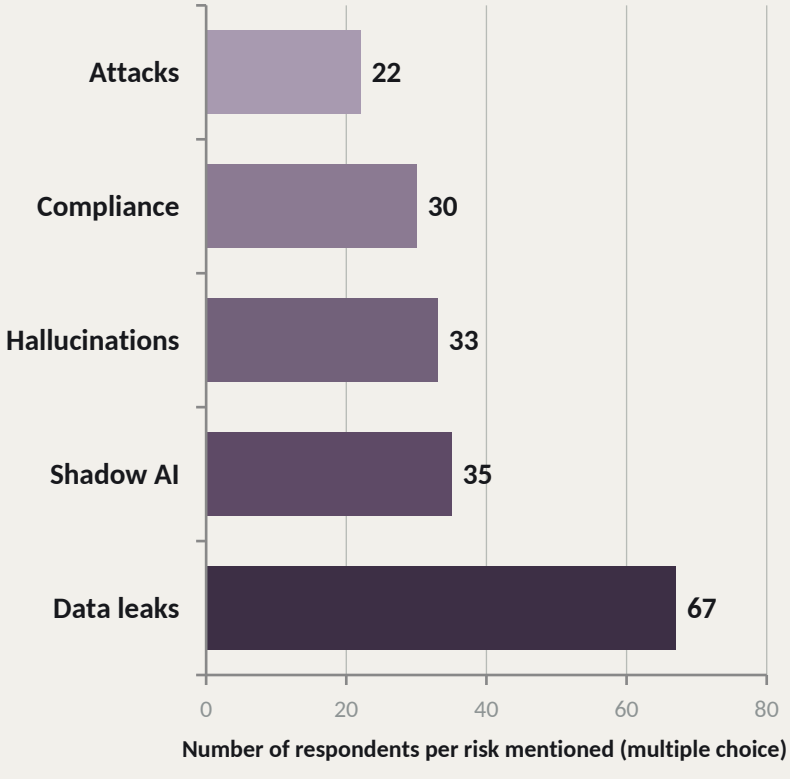

Data leaks dominate the ranking by a wide margin, with 67 mentions, nearly double the second most cited risk.

Three factors explain this prominence. The risk is observable (see the number of data breaches cited earlier), it is already known and covered in classic cybersecurity risk analyses and AI tools are simply a new vector for this existing risk.

At the other end of the spectrum, risk linked to attacks on or via AI tools comes last. The recent launch of Anthropic's Glasswing project strongly challenges this positioning. It reveals that a new model (Mythos, not yet released to the public) has already identified thousands of critical vulnerabilities in major IT systems. The initiative brings together major IT players (e.g., AWS, Google, Microsoft, Cisco) so they can prepare their defenses before the model is publicly released.

1.4 — PROHIBITING DOESN'T MEAN PREVENTING

The conclusion of our discussions on AI-related cybersecurity challenges can be summed up in the title of this section: prohibiting doesn't mean preventing.

Cybersecurity professionals won't be surprised by this. It has been known for years, but in our view, it has never been more true. Companies know that AI is a significant risk factor for their security, they know that AI use exposes their data and that employees are using non-approved tools, but they also know that blocking is often not the most appropriate response.

In a context of strong competitive pressure and a constant fear of missing the turn, banning a tool that generates real productivity gains is seen as counterproductive. Blocking access to AI tools does not eliminate uses, it just makes them invisible. And an invisible use is an uncontrollable use.

In light of this, many of the companies we interviewed have chosen to support these new uses. The challenge is then knowing how, and to what end.

This is where our initial hypothesis ends.

By launching this study, we sought to identify a security issue related to AI, and we succeeded. However, our interviewees led us to take a step back and gain a broader perspective on the issue.

We will explore this shift in the remainder of this study.

01 — GOING FURTHER

Studies complementing our approach

These studies complement and contextualize the data from our approach. We invite you to consult them to go deeper into the topics covered in this section.

ANSSI

Cyberthreat Landscape 2025

France — Critical sectors and economic fabric

This study is the annual snapshot produced by ANSSI on the state of the cyber threat in France. You will find an in-depth analysis of the year's major trends (e.g., erosion of the boundary between state actors and cybercriminals, surge in data exfiltration, targeting of cloud environments, sophisticated social engineering), a focus on AI use by attackers, and useful sector-specific insights to help position your organization.

Read it to understand the cyber landscape onto which AI is being grafted.

Link to the document

KPMG × Les EnthousIAstes

Trends of AI 2026

356 French decision-makers — 47% large corporations, 36% mid-sized companies

The study is conducted by KPMG France and the think tank Les EnthousIAstes. It provides an overview of AI adoption across eight key functions (Marketing, IT, Finance, HR, Customer Relations, Procurement, Supply Chain, Risk & Compliance). You will find a cross-functional reading of the shift from experimentation to deployment and the emergence of new roles such as Chief AI Officers, as well as a focus on the evolving role of the IT department in this new landscape.

Read it to position your organization relative to the dynamics observed in large French companies.

Link to the document

Wavestone

Cyber Benchmark 2025

Large companies — France and international

An annual study measuring the cybersecurity maturity of more than 170 large organizations across 16 themes (e.g., governance, detection, cloud, data, resilience). You will find a global overview of cyber maturity, a sector-specific reading useful for benchmarking (e.g., finance, energy, industry, services, luxury & retail), and above all a section dedicated to securing AI use. It reveals a marked gap between overall cyber maturity and AI-specific maturity, which remains low.

Read it to understand where AI-related cyber blind spots lie and anticipating the work ahead.

Link to the document

CESIN

Baromètre Cybersécurité 2025

CISOs — Large French companies

The CESIN Barometer is the annual reference survey on cybersecurity in French companies, based on responses from 397 CISOs. You will find the 2025 snapshot of the cyberattack landscape, the evolution of the CISO role toward more strategic and cross-functional functions, and a focus on Shadow AI, which in just one year has become the user behavior deemed most risky.

Read it to understand how CISOs perceive and manage new AI-related risks.

Link to the document

02 — AI IMPACT ON OVERALL BUSINESSES

The bigger picture that emerged

2.0 — EIGHT KEY FINDINGS

The interviews opened our eyes to topics we had not previously considered. As the conversations unfolded, we discovered that the arrival of these new tools had impacts well beyond the "simple" question of security. Some impacts appear to be confined to certain companies, while others are shared by all our respondents. They are often framed differently, but we have grouped them into eight takeaways.

These takeaways point toward the same conclusion: companies are facing a major technological (r)evolution that is complex due to its speed and impact, and therefore difficult to grasp and harder to absorb.

Here are the eight takeaways that we will detail in the rest of the document:

01
The first employee-led digital revolution

02
Fear of missing the turn

03
A partial awakening

04
IT teams caught in the crossfire

05
Support functions' expertise challenged

06
Cognitive effects flying under the radar

07
AI moving faster than the organization

08
Systemic exposure

2.1 THE FIRST EMPLOYEE-LED DIGITAL REVOLUTION

The observation shared by all the companies surveyed in our study is that AI did not come in through the front door. It was driven by curious employees, eager to save time, won over by a demonstration or a peer recommendation. Unlike other digital transformations, it was not the result of a senior leader's vision, a dedicated project, or an approved budget. It quietly embedded itself into employees' day-to-day habits.

Usage came before any framework or governance was put in place by the company.

02 — AI IMPACT ON OVERALL BUSINESSES

The bigger picture that emerged

2.1 —USE CASES AHEAD OF THE FRAMEWORK

Concentrated and shallow usage

We were surprised by the reality of actual usage. It is far less spectacular than the promises that dominate the headlines. In 90% of the companies surveyed, the majority of use cases are concentrated within a narrow scope: writing assistance, translation, document summarization, information retrieval, and coding support. AI tools amount to advanced office productivity software or an improved search engine, but do not correspond to any deep transformation of business processes.

Beyond the scope of use, we were also struck by the limited spread of advanced AI usage. It is driven by a handful of curious or particularly tech-savvy individuals, while the majority of employees remain on the sidelines. One respondent put it this way:

> *"We rolled out Gemini licenses for everyone.*
> *[After a few months,] 50% of users had never used their license."*

The gap between deployment and actual adoption limits the impact within companies and therefore the expected return on investment from these tools.

A Three-speed company

Beyond adoption figures, our interviews reveal an organizational divide that is beginning to take shape. One respondent described it this way:

> *"Usage is very people-dependent. It creates a three-speed company: those who go for it, those who go along for the ride, and those who put the brakes on."*

These three speeds coexist within the same organization, sometimes within the same team. They create tensions between colleagues and even between teams, and a widening productivity gap between those who have mastered these tools and those who have yet to use them.

Grey AI, a new path for AI usage

Shadow AI is a well-recognized phenomenon among our respondents. 86% believe it exists within their organization. However, our interviews led us to discover a related concept, sitting at the boundary between Shadow AI and legitimate use. A CISO from the healthcare sector brought it to our attention and named it: Grey AI.

> *"Grey AI is AI that has been built into legacy tools that are already authorized within the company, but whose existence we are unaware of."*

Many established software vendors have integrated AI features into their products. For SaaS (Software as a Service) solutions, users see these features appear overnight. They are deployed without anyone in the company having assessed what happens to the data, which models are being used, or which servers the prompts are sent to. The tool itself remains authorized, but should its new AI layer be as well?

Grey AI is harder to detect than Shadow AI because it hides behind the legitimacy of the tool. Users, reassured by the familiar and routine nature of the tool, feed it data they would not naturally entrust to a dedicated AI tool.

2.2 — FEAR OF MISSING THE TURN

Almost all of our respondents describe a widespread and constant pressure. It comes from both inside the organization (i.e., leadership, business units, employees) and outside (e.g., vendors, suppliers, consultants, media, influencers, conference speakers).

Pressure built from the outside

The AI market has broadly been built around a dual narrative. On one side, promises of significant gains: productivity increases, full process automation, competitive advantage for early adopters, and the list goes on. On the other, a threat: those who fail to adopt AI will quickly fall behind.

A large part of the AI ecosystem (e.g., vendors, consultants, influencers, specialist media) amplifies this narrative because their ability to be heard often depends on outbidding one another and stoking a sense of urgency. One respondent summed up the mood with this phrase:

> *"The constant hype on LinkedIn gets users dreaming. They think AI is going to solve all our problems, but IT isn't magic."*

From our perspective, the communication itself is not the issue. The problem is the recurring gap between the promises on display and the reality of the current implementation.

Internal pressure from business units

Within organizations, this narrative crystallizes in the form of urgent, poorly scoped requests from business teams. IT departments face employees and managers who have seen a demo or read an article that convinced them of the need to move immediately and deploy AI tools without delay. A recently hired AI project manager shared this:

> *"The sky-high promises around AI make people lose their rational thinking. When you dig into the actual need, AI isn't even always the best solution."*

This pressure has a name: FOMO — Fear Of Missing Out. The fear to miss the turn.

Accelerating adoption does not create value

The FOMO phenomenon has been well documented for some time, particularly in the context of social media. On the AI side, an academic study published in 2025 in the journal Technology in Society documents the concept of AI FOMO. It shows that employees exposed to intense AI discourse develop a specific form of anxiety linked to the fear of being left behind or made obsolete, regardless of their actual ability to use these tools.

The effect is counterproductive: FOMO accelerates adoption decisions but risks undermining adoption over the long term. Projects launched under pressure and without a defined use case fail more often.

95%

of business AI initiatives show no return on investment

MIT - THE GENAI DIVIDE – STATE OF AI IN BUSINESS 2025

2.3 — A PARTIAL AWAKENING

In our view, it would be unfounded to say that business leaders are ignoring AI. Throughout our study, we observed that awareness is genuine and that the topic is being taken seriously. However, there is a gap between being aware of a challenge and being able to drive an appropriate response, and that gap is what our interviews brought to light.

Just another wave, not a break from the past

There have been previous strategic shifts that have properly taken root, like digitalization, the Cloud, or SaaS. On the other hand, there are those others that ultimately amounted to little more than announcements with limited real-world impact, such as NFTs, the metaverse, or enterprise blockchain.

Every wave, real or anticipated, has generated its share of promises, urgency, external consultants, and pilot projects. Once the dust settled, even when the technology delivered on its promises, it ended up being absorbed into the existing landscape without much fundamental structural rethinking.

Our interviews suggest that several leaders are approaching AI through the same lens: an opportunity to seize, risks to manage, but not necessarily a reason to fundamentally rethink their organization. This approach is understandable. Leaders must constantly weigh dozens of strategic priorities and cannot restructure their organization on a daily basis.

58%

of executives consider AI a medium-term survival issue

BPI FRANCE - LE LAB SUR L'IA 2025

Optimizing rather than transforming

We looked for a study to challenge the findings from our interviews. The only French study we found that addresses this topic is the one by BPI. However, the data was collected at the end of 2024, an eternity in the world of AI. We nonetheless found data points that extend our analysis.

Only 32% of SMEs and mid-sized companies were using AI, and a majority of leaders had yet to formalize a strategy. Among those using AI, 54% started with free tools. The stated priority was optimizing what already exists: cutting costs or improving performance. These findings could be indicators of the genuine but partial awareness we describe in this section, though a great deal may have changed since then.

AI amplifies, it does not fix

The people we interviewed pointed to a blind spot that sometimes exists in their leaders' thinking: AI does not behave like previous technology waves.

The cloud brought elasticity. SaaS offloaded infrastructure. Those waves integrated into what was already there. But AI works differently. It plugs directly into a company's data, processes, and capabilities, and amplifies what it finds.

Where things work well, that amplification will mostly produce real gains. But where things do not work (e.g., silos, technical debt, poorly defined processes) it will produce the opposite effect. It will potentially make dysfunctions faster, more visible, and more damaging.

This amplifying factor is what sets AI apart from previous waves.

2.4 — IT TEAMS CAUGHT IN THE CROSSFIRE

We now turn to dynamics that are specific to IT departments as they navigate the arrival of AI tools. To start, let us consider the primary mission of IT departments: they must ensure that systems run reliably over the long term. They are organized accordingly: evaluation, validation, implementation, control, maintenance. Starting from this premise, we identified two phenomena that are disrupting and complicating that mission.

Vibe coding

Before going further, a definition: vibe coding is the ability to generate functional code through natural language prompts, without necessarily understanding what the code actually does. The immediate result is appealing to non-developers. They become capable of independently delivering a working micro-application within a matter of hours. The medium-term result, however, is a source of concern for our respondents:

> *"Vibe coding is a generator of shadow IT and technical debt in the medium term, and it's going to become unmanageable."*

This ambivalence is significant. Vibe coding occupies a particular place in our analysis because it divided our respondents. For some, it is seen as an "empowering" tool for non-developers. Others warn of the medium-term risks. From their perspective, these micro-tools spread through organizations because they address a genuine business need, often with efficiency and immediacy. They are nonetheless problematic because no one fully knows how to maintain, audit, or assess the security implications of what has been built. And if something goes wrong, it will be difficult for IT to step in.

This dynamic is the first manifestation of what we call in this report the illusion of competence: a sense of mastery that masks a growing dependency on tools whose inner workings and limitations are not fully understood. We will return to this in section 2.6.

From our respondents' perspective, the challenge is not to ban, but to provide guardrails without stifling, so that those practices align with the company's objectives and integrate into the existing IT landscape.

Agentic AI: the second wave

Everything covered so far relates to the first AI wave: that of generative tools, assistants, and vibe coding. A second wave is appearing on the horizon, and our respondents view it with even greater wariness: agentic AI.

AI agents are systems capable of acting autonomously (browsing the web, executing code, modifying files, interacting with APIs, and chaining tasks together) without human oversight. Several respondents expressed explicit concern:

> *"The side effects with agentic AI are going to be multiplied many times over."*
> *"We're going to have to find the right balance between human oversight and autonomy."*

If the first wave put IT teams under pressure, the second risks overwhelming them entirely. The terrain is even more uncertain: autonomous actions, cascading errors, diffuse accountability, and more.

Agentic AI has not yet been widely deployed in French SMEs and mid-sized companies, but it is coming. Organizations have not yet fully mastered the first wave, and they must already begin preparing for the second, which looks set to be even more structurally significant.

2.5 — SUPPORT FUNCTIONS' EXPERTISE CHALLENGED

This section is based on a smaller number of respondents, but the implications seem significant enough to warrant dedicated coverage. AI appears to be shifting another balance within organizations, one less visible than Shadow AI or technical debt, yet potentially more structurally consequential in the long run. Support functions are seeing their position change.

A long-standing model in motion

For decades, support functions (e.g., HR, legal, finance, communications) have occupied a distinctive position within companies. They held specialized knowledge that other teams did not have. Those teams therefore depended on support functions to access it, which allowed support functions to set their own timelines, formats, and priorities, and more broadly, their own ways of working. This power dynamic was unspoken but real. AI tools are now cracking this established model. An employee can now draft a standard contract, prepare a financial analysis, produce an HR document, or get a first-level legal answer without consulting internal experts. The answer is often approximate, sometimes outright wrong, but it is immediate. What is more, it appears convincing enough that the employee is satisfied with it. One respondent shared this:

> *"We've already had to deal with a team that signed a contract with a supplier without going through procurement or legal, something we'd never seen before!"*

This kind of situation may no longer be an exception in the years ahead. A study published in January 2026 by the Collective Intelligence Project, conducted among 6,000 people across 70 countries, found that 65% of respondents trust AI chatbots to act in their interest, ranking them in the top three, just behind their family doctor and public research institutions.

Decisions that only appear informed

The main risk is not that support teams will lose their jobs. The real risk, in our view, is that employees will be tempted to make significant decisions (e.g., contractual, regulatory, financial) without consulting internal experts, relying solely on tools that have limited knowledge of the company's context and bear no responsibility for the consequences. One respondent shared this observation:

> *"The problem with AI tool outputs is that they always seem plausible, but they are not always accurate."*

This is the same illusion of competence we identified in section 2.4, but this time applied to expert knowledge. The employee believes they can bypass the lawyer, the finance specialist, or the HR manager because the tool gives them an immediate and plausible-sounding answer. But an answer is not a reasoned judgment.

An organization that leaves its employees to navigate alone, without guardrails or guidance from expert functions, exposes itself to errors with serious legal, financial, and reputational consequences.

A transition to be managed

In our view, this change is not inherently negative. It is an invitation for support functions to reinvent themselves, to move from the role of knowledge gatekeeper to that of architect and enabler. If this dynamic becomes widespread, support functions will need to recognize it and adapt in order to continue fulfilling their role fully.

2.6 — COGNITIVE EFFECTS FLYING UNDER THE RADAR

We have already touched twice on manifestations of cognitive effects that AI tools produce in employees. It is difficult for us to provide a detailed analysis of all the biases activated by AI use. The technology is recent and we lack the hindsight to understand its full, long-term impact. Nevertheless, we have gathered a number of findings that have already been identified by research. This topic strikes us as important because it goes to the heart of our relationship with our own competence and our ability to assess the quality of what AI produces.

> *"Everyone feels like they are an expert: AI gives everyone who uses it a sense of competence on every subject."*

The mechanism our respondent describes is well documented, and its implications extend far beyond the workplace. Generative AI consistently produces outputs that appear plausible and high-quality. Often, the result is genuinely better than what the user would have produced alone. However, this improvement comes with a cognitive cost that is easy to overlook: the user delegates without learning.

The reverse Dunning-Kruger effect

Since 1999, cognitive psychology has described the Dunning-Kruger effect: individuals with the least competence in a given area tend to overestimate their ability, while experts, aware of the complexity involved, tend toward humility. Generative AI produces the opposite effect.

A study published in 2026 in Computers in Human Behavior by researchers at Aalto University (Finland) shows that using ChatGPT eliminates the classic Dunning-Kruger effect. All users, regardless of their actual level of competence, overestimate their performance when using AI. Those who consider themselves most comfortable with AI are also those who overestimate themselves the most. In other words, the more we use AI, the less capable we are of lucidly evaluating what it produces and our own level of ability.

Cognitive offloading

The same study identifies a second mechanism: cognitive offloading, or the act of delegating one's thinking to the tool without verification or critical reflection. The majority of participants were satisfied with a single prompt and accepted the response as given.

The long-term consequences are still poorly understood, but we can formulate the following hypothesis: a skill we no longer practice will atrophy. The developer who simply validates generated code may gradually lose the ability to write and audit it. The lawyer who delegates research to an LLM may lose the depth of their reasoning. The list could extend across many professions. An immediate productivity gain could, in short, conceal a slow erosion of collective competence.

What this means for organizations

We have limited hindsight on the implications and the solutions to put in place. Our conversations allowed us to identify some initial avenues for action that need to be confirmed and built upon:

- Training people to critically evaluate AI-generated outputs, as much as training them to use AI tools themselves
- Maintaining spaces for practice without AI in order to preserve foundational skills
- Distinguishing between mastery of an AI tool and mastery of a field of expertise

2.7 — AI MOVING FASTER THAN THE ORGANIZATION

Our respondents described a feeling of constant tension. When we probed further, we came to understand that it stemmed from the coexistence of two incompatible rhythms: that of AI and that of their organizations.

Two worlds, two speeds

The AI ecosystem is evolving at a pace unprecedented in the history of enterprise technology. Vendors release continuous updates, launch new models, add features, and integrate ever more deeply with the rest of the IT stack. The capabilities of these tools grow from one quarter to the next at a rate that shows no sign of slowing. One respondent put it well:

> *"The AI ecosystem moves very fast, too fast. I actively follow the space and yet I feel like I'm falling behind."*

Companies operate on a much longer timescale. A major IT project takes 9 to 18 months. From identifying a need, analyzing it properly, finding the right tool, integrating it into the existing systems, and training users. These steps are necessary, and most of them cannot be compressed. Every project manager knows that speed often comes at the expense of quality or budget.

The problem is not that organizations are slow, but that the speed of AI makes that slowness visible, and therefore harder to accept for those driving projects.

Decisions made on already obsolete tools

This mismatch has a direct consequence: by the time an organization has finished evaluating, validating, and deploying an AI tool, that tool has often already been superseded by a new version or a new competitor.

IT and business teams find themselves facing a difficult choice:

- wait for a stable picture, at the risk of falling behind in ways that will be hard to recover from,
- or decide quickly and commit to choices that will be obsolete before they are even fully deployed.

Whichever path is chosen, the exercise becomes increasingly exhausting over time. Every week, teams must evaluate, prioritize, and make judgment calls on requests and announcements under tight time constraints. Decision fatigue sets in gradually.

What research confirms

The timing mismatch is not unique to AI, but AI has intensified it. Work in change management consistently points to the same conclusion: a successful technology transformation requires a minimum cycle of 18 to 24 months to become embedded in practice. The AI ecosystem renews its core capabilities every 3 to 6 months.

In our view, the answer to this mismatch does not appear to be moving faster, but discerning better. Companies can learn to distinguish what merits immediate adoption, what can wait, and what amounts to noise. Easier said than done.

2.8 — SYSTEMIC EXPOSURE

The seven previous takeaways are not independent problems. They are symptoms of the same phenomenon we mentioned in the context portion of this study: the gap between the speed of AI and organizations' capacity to absorb it. Taken together, they feed into one another and produce a systemic exposure that is difficult to measure.

A compounding effect

The takeaways do not simply add up, they reinforce each other. Poorly controlled usage fuels pressure from business units. That pressure reaches leaders who have not yet formalized their strategy, and who sign off on rushed projects. Those projects then land on IT teams that are already stretched, creating technical debt that will become increasingly difficult to manage. Meanwhile, support functions lose their role as a safeguard, users overestimate their own competence, and none of this keeps pace with the evolution of the technology itself. The real risk is a gradual loss of visibility, and control over what AI is actually doing within the organization and what value it is actually generating.

The pervasive risk of data leakage

Data leakage is the most immediate risk. Our respondents place it at the top of their AI-related concerns. The at-risk situations are numerous: customer data sent to consumer-facing LLMs, strategic information in prompts within existing tools, personal data processed without a legal basis, source code shared with third-party models, and more.

This risk is all the more difficult to manage because it is diffuse. It does not stem from an external attack, but from the everyday (and often well-intentioned) behavior of employees. It accumulates without any visible incident, until it becomes a regulatory or competitive problem.

N°1

Data leakage is the number one risk identified by respondents to our survey.

The compliance risk

NIS2, DORA, the AI Act. The European regulatory landscape is becoming increasingly dense. These frameworks impose specific obligations around AI system governance, usage documentation, risk management and the traceability of algorithmic decisions.

Organizations that have not mapped their AI usage will struggle to demonstrate compliance. Yet usage is multiplying faster than governance can be established. Every month that passes makes that mapping harder to reconstruct.

The strategic risk

Beyond the immediate and regulatory risks lies a longer-term risk. As we noted in section 2.3, these new tools act as an amplifier of what already exists. They reinforce an organization's strengths just as much as its weaknesses. AI becomes a resonance chamber for how the company actually operates.

The entire challenge of implementation will be to amplify strengths without amplifying weaknesses. In a market where promises are commonplace, it will be very tempting to ask AI to compensate for a weakness, but what does not work before AI rarely works better after it.

Over time, the competitiveness gap risks widening between organizations that know how to build on their strengths when implementing AI, and those that do not.

2.9 — A DIAGNOSTIC FOR ACTION

At the beginning of this study, we had a focused hypothesis: AI is a major cybersecurity challenge for companies. Through our conversations, that hypothesis proved to be both correct and insufficient. Correct, because the cyber risks are real and intensifying. Insufficient, because the issue is not purely one of security. It runs deeper and can be summed up in a single sentence: organizations are structurally unprepared to absorb a technology that evolves faster than they can adapt.

This gap between the speed of AI and organizations' capacity to absorb it is the thread running through everything we have written. It takes different forms depending on the angle, but it is always the same underlying phenomenon.

What this study offers and what it does not

Let us be transparent about the nature of our contribution: we have not made any revolutionary discoveries or documented previously unknown phenomena. Shadow AI, pressure from business units, technical debt, the illusion of competence, the timing mismatch, these realities are described in other studies, experienced across all organizations, and are likely already familiar to the reader who has followed us this far.

What we are offering is more modest. We wanted to consolidate an overall picture grounded in fieldwork. We interviewed IT decision-makers and professionals, cross-referenced their input with academic and institutional research, and attempted to organize these takeaways into a framework that is useful to those who must make decisions.

We also wanted to shine a light on a few mechanisms that we feel are underrepresented in current discourse: Grey AI, the illusion of competence, the amplification effect, and AI FOMO. If these concepts help a business leader or CIO put words to what they are observing in their organization, we will have achieved our goal.

Clarity without paralysis

The picture we have painted could be read as a worrying diagnosis, but that is not the message to take away, as we see things differently.

The gap between the speed of AI and organizations' capacity to absorb it is only a starting point. Recognizing that gap, understanding its mechanics, and accepting it strikes us as the first step toward addressing it effectively.

The final section that follows offers concrete courses of action, also drawn from our conversations and our reading. They are not universal, and they are not a substitute for the knowledge each organization has of its own context.

We firmly believe that there is a great deal to build, and that the companies that manage to find their own transformation tempo will be the ones that extract the most value from it.

02 — GOING FURTHER

Studies complementing our approach

These studies complement and contextualize the data from our approach. We invite you to consult them to go deeper into the topics covered in this section.

Aalto University

Reversed Dunning-Kruger and AI

Academic study — published February 2026

This Aalto University study highlights two cognitive mechanisms activated by AI use. You will find a demonstration of the reversal of the Dunning-Kruger effect in the presence of AI and an analysis of cognitive offloading.

Read it to understand the cognitive effects of AI on its users.

Link to the document

ScienceDirect × Technology in Society

AI FOMO

Academic study — published 2025

This academic study formalizes the concept of FOMO applied to AI in the professional world. You will find a demonstration that AI FOMO is a distinct phenomenon from traditional FOMO, as well as an analysis of the factors that amplify or attenuate it.

Read it to understand how AI activates specific psychological mechanisms in the workplace.

Link to the document

MIT Project NANDA

State of AI in Business 2025

52 organizations, 153 senior decision-makers — 2025

This MIT study analyzes why enterprise AI projects generate no measurable return even when tools are massively adopted. You will also find insights on widely observed Shadow AI in businesses and the factors that distinguish the rare projects that create value.

Read it to understand what tips a pilot into real value.

Link to the document

McKinsey & Company

The State of AI 2025

1,993 respondents — 105 countries

McKinsey's State of AI is one of the broadest international snapshots of enterprise AI adoption. In it you will find an analysis of the "scaling gap" between widespread adoption and actual scale-up, a focus on the emergence of agentic AI across business functions, and the practices that set apart the rare organizations that are generating measurable value from AI.

Read it to identify the levers for scaling up.

Link to the document

Stanford University

Vibe coding & security

47 participants — 2023

This Stanford study measured the effect of AI assistants on the quality of code produced by developers. You will find a demonstration that AI-assisted developers produce less secure code than those coding alone, while being more confident in the quality of their work.

Read it to understand the concrete risks AI poses to code security in the enterprise.

Link to the document

Bpifrance Le Lab

AI in french SMEs

1,209 SME/mid-size company executives — France

This BPI study is the most comprehensive picture of AI adoption in French SMEs. You will find a mapping of AI maturity by sector and leadership profiles, as well as five concrete use cases.

Read it to position your organization and identifying your next priorities.

Link to document

03 — ACTION TRACKS

After the assessment, how do we act?

3.0 — DO WE HAVE THE MAGIC RECIPE?

At the risk of surprising you: no, we do not have a foolproof methodology for deploying AI tools in your organization. Our interviews and research have, however, allowed us to identify some directions and footholds that we believe are good starting points. These are not a substitute for the knowledge each organization has of its own context. We therefore invite you to adapt them to your own vision and operational constraints.

We came across many methodologies built around structured sequences: step 1, then step 2, then step 3. While useful, the linear methodologies we reviewed reach their limits when faced with a technology that evolves faster than transformation cycles. AI governance will need to be more flexible, able to skip steps when speed is needed, or to backtrack when necessary.

3.1 — KNOW WHERE YOU STAND BEFORE YOU ACT

Faced with the pace mismatch, the natural reflex would be to accelerate: launch a project, roll out licenses, train teams. Internal and external pressures push in that direction, as mentioned in part 2.

We are convinced this would be counterproductive. Our interviews and the literature we reviewed confirmed our suspicions. Those who navigate this best are those who take the time to look at what is already happening within their organization before formulating a response. Three angles strike us as particularly useful for this exercise.

Know the real usage

Real usage hides in the Shadow AI documented in part 1, in the Grey AI we described in section 2.1, and in micro-applications developed through vibe coding by resourceful employees. An organization that relies solely on its official AI tool inventory is building its strategy on a flawed picture. Mapping real usage is not a one-off audit. Practices shift quickly. This exercise needs to be revisited regularly.

Assess data maturity

AI tools create value from the data they are given. Companies rarely lack data; the complexity usually lies in the ability to extract and use it. An organization with well-organized, accessible data has a genuine advantage. The methodologies we reviewed typically apply three criteria: does the data exist? is it accessible? is it of sufficient quality? Without usable data, AI pilot projects will stall or never scale.

Leadership posture

The third angle is less technical: it concerns the posture of leaders toward AI, which shapes all the decisions that follow. We described in section 2.3 the risk of approaching AI as just another wave. This view is understandable, and yet it limits an organization's ability to fully integrate AI.

Several questions are worth asking before structuring an approach. Is AI being driven by the executive leadership or by the IT department? Has the organization formalized an ambition, transformation, optimization, or cautious status quo? How will employees and managers be involved in the process?

The answers will vary from one organization to the next. Asking the questions provides the compass with which to navigate.

3.2 — GOVERNANCE THAT KEEPS PACE

Once the lay of the land is known, the question of implementation arises: how do you build governance that is neither stifling nor absent?

Classic governance models are designed for long cycles, steering committees, validation procedures, and so on. These mechanisms have proven their worth on transformations that followed a predictable rhythm, but they struggle to keep up with AI.

Guiding rather than gatekeeping

We believe the shift in governance lies in the approach. Effective AI governance is governance that guides business teams and informs decision-making. It gives teams the means to decide quickly and well, while maintaining a shared direction across the organization.

In practice, this could mean establishing common principles rather than procedures. For example: clear criteria for distinguishing high-stakes AI use cases, explicit accountability by type of decision, thresholds that trigger a project review, proof of value rather than proof of concept, and so on. The challenge will be striking the right balance between freedom of initiative and the organization's overall framework.

Building collective momentum

The response to these challenges will not come from any one department working in isolation. Not IT alone, not business units alone, not leadership alone. The AI turning point must be seen as a cross-functional company-wide endeavor that brings everyone to the table.

The difficulty will also come from a difference in perspective. Leadership seeks efficiency, performance, and strategic alignment. Employees look for what concretely helps them in their day-to-day work, like tools that works, a practice that fits naturally, or a meaningful improvement in comfort. These two perspectives are not opposed, but they do not naturally find each other.

Several of the most mature organizations we interviewed offered a similar response. They created dedicated mixed teams focused on AI, sometimes called an AI Factory. These teams bring together business, IT, and data profiles around concrete use cases. They allow ground-level needs to surface, technical and business expertise to intersect, and projects to be framed within a shared reference point. The strategy has its limitations and isn't one size fits all, but the underlying principle remains compelling.

Compliance as a lever

Whenever governance comes up, regulatory considerations emerge almost immediately. We sometimes sensed a degree of weariness among our respondents when faced with the accumulating layers of regulation and standards they must navigate: GDPR, NIS2, DORA, the AI Act, and more.

The obligations imposed by the AI Act, taken in isolation, can seem burdensome: mapping usage, tracing algorithmic decisions, ensuring human oversight, documenting training data. However, these obligations largely overlap with what sound AI governance should be doing anyway. Mapping usage is useful for steering based on real data. Tracing decisions is useful for auditing in the event of an incident. Clarifying human oversight is useful for assigning accountability.

The AI Act can then be seen as yet another constraint or as a framework to help structure AI governance.

3.3 — TRAIN AND BRING PEOPLE ON BOARD

Deploying an AI tool will be the "easy and fast" part of the project. As with any transformation, the real challenge lies in adoption: getting employees to take ownership of the tool and accept it. Boston Consulting Group pitches a 10-20-70 success ratio: 10% of the effort should focus on algorithms, 20% on data and infrastructure, and 70% on people. Without being fans of sweeping formulas, this breakdown clearly illustrates the importance of training and change management.

Training to use and to evaluate

We documented in section 2.6 certain biases associated with AI use. We do not yet have enough hindsight to fully understand the impact of these biases in the workplace. However, it appears that training employees on how to operate these tools correctly will not be enough. Companies will also need to train employees on how to collaborate effectively with AI tools in order to avoid medium-term issues, over-reliance on AI, loss of skills, difficulty evaluating what AI produces, and more. This can be achieved through many existing mechanisms (e.g., awareness campaigns, traditional training, micro-learning) or through approaches yet to be invented, for example, a dedicated practice space without AI.

Bringing all the three speeds onboard

We described in section 2.1 the three-speed company: those who go for it, those who go along for the ride, and those who put the brakes on. A uniform training system will quickly reach its limits and will not address this divide. The fast movers will be bored, the resisters will disengage, and the middle group risks being left behind.

Several approaches were described by our respondents to narrow this divide:

- Leaning on "ambassadors": identifying early AI adopters within teams, turning them into operational relays for their colleagues, and gradually widening the circle.
- Leaning on support functions: they know the practices, constraints, and friction points of the teams they work with. They can take on an advisory and relay role in spreading an AI culture.
- Implementing contextual learning systems: using tools or methods that allow everyone to learn at their own pace and according to their day-to-day needs.

Each organization will need to find the combination of approaches that works for it.

3.4 — KEY TAKEAWAYS

If we look at the major technology waves that came before, each one divided organizations into three groups: those that missed the turn, those that managed the transition, and those that seized the opportunity to reinvent themselves. AI could follow the same pattern with one key difference: it gives organizations less time to choose their direction.

From the conversations and reading we have shared, we would like to present three recommendations:

- Gaining a clear picture of what is actually happening within your organization.
- Rethinking governance so that it keeps pace with change rather than being overwhelmed by it.
- Training employees both to use these tools and to evaluate what they produce.

These three angles do not follow a fixed sequence. They should be seen as three mutually reinforcing priorities.

Three questions to look ahead

Before closing this document, we offer a brief forward-looking exercise. Try asking these three questions within your organization:

- What tasks exist today because information is scarce or difficult to process, rather than because they create value?
- If your teams were freed from 40% of their operational workload, what could they propose and build?
- Which processes or actions currently carried out by your teams should remain human work, even if a tool could do it faster?

The first two explore what could change. The third defines what should not.

Acknowledgement

This report would not have been possible without the 135 people who agreed to complete our survey or take the time to speak with us in interviews. We sincerely thank them for the quality and candor of their contributions.

We also thank you, the reader, for staying with us to the end.

This document is just one step for us. If these findings resonate with your context, or if you would like to challenge them, we would be happy to continue the conversation, on LinkedIn, by email, or through whatever channel works best for you.

03 — GOING FURTHER

Studies complementing our approach

These studies complement and contextualize the data from our approach. We invite you to consult them to go deeper into the topics covered in this section.

Boston Consulting Group

Deploy, Reshape, Invent

Consulting firm — AI methodological framework

BCG proposes a strategic framework structured around three value-creation mechanisms (i.e., Deploy, Reshape, Invent) and a counterintuitive investment ratio: 10% on algorithms, 20% on data and infrastructure, 70% on people and processes. You will find a reading of the budget trade-offs in successful AI transformations, as well as quantitative data on company trajectories according to their maturity level.

Read it to revisit budget allocations in light of a field-tested ratio.

Link to the document

HEC Montréal

AI Implementation Methodology in Organizations

Academic institution — Methodological framework

HEC Montréal has structured a five-step approach to AI implementation in organizations, with particular attention to data quality, accessibility, and existence. You will find an iterative approach that accounts for back-and-forth between steps, as well as concrete criteria for assessing data maturity before launching an AI project.

Read it to validate the feasibility of an AI project starting from the data.

Link to the document

Ministère du Travail × Inria

Deploying AI at Work

French public program — Operational guide

Labor IA, a partnership between the French Ministry of Labour and Inria, has published a nine-step guide to help SME and mid-sized company leaders deploy AI. You will find an approach centered on social and technological dialogue, as well as specific attention to algorithmic governance and transparency toward employees.

Read it to structure an AI deployment within a French framework.

Link to the document

Deloitte

State of AI in the Enterprise

Consulting firm — Study on AI adoption

Deloitte's annual study measures the state of AI adoption in enterprises and identifies the factors that distinguish successful transformations from others. You will find a demonstration that organizations that invest in change management are 1.6 times more likely to see their AI initiatives exceed expectations, as well as an analysis of employee onboarding levers.

Read it to understand why change management matters more than technology.

Link to the document

APPENDIX

2026 Synthesis

Detailed survey results

95 survey respondents : CIOs, CISOs, and IT decision-makers in France and Belgium

DETAILED SURVEY RESULTS

01 – What is your primary role?

Single-choice question

Answer	Share
CISO / Information Security Manager	**50.0 %**
CIO / CTO / IT Manager	**43.1 %**
Other IT Managers	**5.6 %**
DPO / Legal Counsel / Compliance Officer	**1.3 %**

02 – What is the size of your organization?

Single-choice question

Answer	Share
SME (< 250 employees)	**41.7 %**
Mid-sized company (250 to 5,000 employees)	**38.9 %**
Public sector organization (local authority, hospital group, university hospital, etc.)	**12.5 %**
Large enterprise (> 5,000 employees)	**6.9 %**

03 – What is your company’s policy regarding AI tools?

Single-choice question

Answer	Share
Restricted: only an "Enterprise" version (e.g. Copilot) is authorized	**40.3 %**
Controlled: access permitted, governed by a policy or mandatory training	**36.1 %**
Open: unrestricted access, no particular limitations	**11.1 %**
Under consideration: no official position has yet been established	**11.1 %**
Prohibited: complete technical block	**1.4 %**

DETAILED SURVEY RESULTS

04 – Do you believe Shadow AI exists within your organization?

Single-choice question

Answer	Share
Yes, significantly: common practice despite existing rules	**43.8 %**
Yes, but marginal: a few isolated users	**42.2 %**
No, it is under control: users comply with the policy	**7.8 %**
Hard to estimate: no real visibility into these flows	**6.2 %**

05 – How do you currently manage the risk of data leakage to AI tools?

Single-choice question

Answer	Share
No technical measures — awareness and accountability only	**56.9 %**
Classic DLP tools — our current tools filter some outgoing data	**26.4 %**
Network blocking — traffic blocked via proxy / firewall	**16.7 %**

06 – Which AI-related risk is most critical to you today?

Multiple-choice question (up to three answers)

Answer	Nb of mentions
Data leakage: sending confidential or personal data to AI models	**67**
Shadow AI: not knowing who is using what	**65**
Hallucinations: factual errors affecting the relevance of outputs	**33**
Compliance: GDPR / AI Act exposure	**30**
Attacks: prompt injection, AI jailbreaking	**22**

07 – What is your biggest obstacle to deploying a dedicated security solution?

Multiple-choice question (up to three answers)

Answer	Nb of mentions
Budget: no allocated budget at this stage	**45**
Lack of offer: no convincing solution available on the market	**25**
Not a priority: the risk is considered acceptable	**17**
Complexity: concern about degrading user experience (latency, friction)	**16**
Wait-and-see: waiting for Microsoft / Google to integrate these controls natively	**4**

08 – Which technical approach do you favor?

Single-choice question

Answer	Share
Gateway / Proxy: invisible layer filtering traffic to AI tools	**41.7 %**
Dedicated platform: internal chat interface (a "home-built ChatGPT")	**29.2 %**
Browser extension: plugin installed on the client device	**25.0 %**
API: invisible component integrated into internal applications	**4.1 %**

09 - What features would be essential for you to purchase this type of product?

Multiple-choice question (up to three answers)

Answer	Nb of mentions
Audit / logs: knowing who sent what	**50**
Pseudonymisation: replacing sensitive data with aliases	**45**
Business contextualization: smart blocking of sensitive concepts (e.g. "strategy")	**42**
Security: blocking malicious code or inappropriate content	**42**
Reversibility: the AI responds with aliases, the tool restores the real data	**16**

10 - How urgent is it for your organization to adopt such a solution?

Single-choice question

Answer	Share
Exploratory: monitoring underway, potential acquisition within 12 to 24 months	**61.1 %**
Low: not a priority for the coming year	**36.1 %**
Immediate: actively looking for a solution (2026 budget approved)	**2.8 %**

The path is narrow. The stakes are high. But with the right visibility, every step is deliberate.

TERRIA CONSEIL

contact@terriaconseil.fr

@Anthony Poyet @Eleonore Wiggins @Julien Kilo