AIMONGER
WHITE PAPER

AI Discovery

The False AI Race: Strategic Patience for Leadership Teams

Much of the public AI narrative is framed as a race with a single finish line. That frame produces panic buying, shallow pilots, and staff anxiety. Eurostat shows 20.0 per cent of EU enterprises with 10+ employees used AI in 2025, while McKinsey shows enterprise EBIT impact remains narrow among AI-using organisations. The useful posture is strategic patience with operating urgency: absorb a few workflows deeply rather than perform speed.

David Saliba . 2026-07-01 . aimonger.com/whitepapers/strategic-patience-false-ai-race/

The problem in one sentence

Feeling behind is not a KPI. Boards need a tempo narrative: fewer workflows absorbed deeply with metrics, not a panic portfolio that funds demos to quiet headline anxiety.

Strategic patience, plainly: fund a short list of workflows with owners, ninety-day gates, and kill criteria, while moving fast on data access and operator training. **In practice:** three production workflows with cycle-time metrics beat fifteen pilots that never graduate.

False AI race, in short: the belief that every model release and hyperscaler headline creates a single finish line your firm must cross this quarter, regardless of absorption capacity. **Worked case:** OpEx rises with Copilot seats while EBIT attribution stays below 5 per cent because no workflow was redesigned.

Panic is a procurement strategy

When leaders feel behind, they buy tools. Buying tools creates the feeling of catching up. It rarely creates absorption.

In the boardroom, the feeling of being behind is usually imported from elsewhere: US hyperscaler spend, startup launch cycles, viral demos, and vendor decks that turn every model release into a deadline. That global story is real at the infrastructure layer. It is not a useful operating plan for a mid-market firm in Malta, Milan, Munich, or Madrid.

Gartner expected GenAI APIs and GenAI-enabled applications to be in mainstream enterprise use by 2026. That mainstreaming increases noise. McKinsey's impact figures show noise is not value. The false race says "more." The operating race says "deeper."

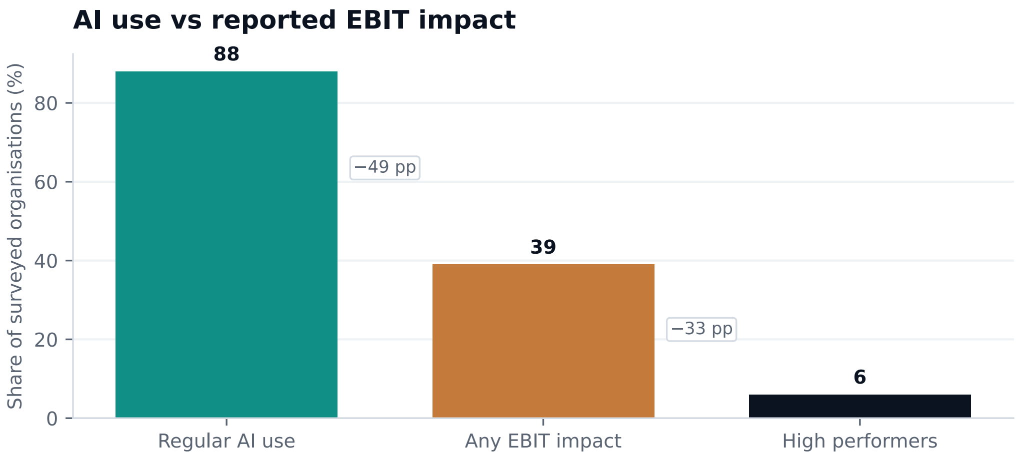

McKinsey State of AI 2025 (survey). AIMonger redraw.

Figure 1. Organisations regularly using AI vs those reporting enterprise EBIT impact vs high performers (>5% EBIT). Plain read: most firms have AI somewhere; few can point to profit. **On the P&L:** Copilot seats show up in OpEx; the P&L still looks like last year unless workflows changed. Source: McKinsey State of AI 2025 (survey). AIMonger redraw.

McKinsey's 2025 survey reports nearly nine in ten respondents say their organisations regularly use AI in at least one function, while only 39 per cent attribute any enterprise-level EBIT impact to AI, and about 6 per cent qualify as high performers attributing more than 5 per cent of EBIT to AI. High performers redesign workflows and scale what works; they do not win by licence count.

Eurostat reports 20.0 per cent of EU enterprises with 10 or more employees used AI in 2025, up from 13.5 per cent in 2024. Large enterprises reached 55.0 per cent. Many peers are still early; "behind the world" is often false.

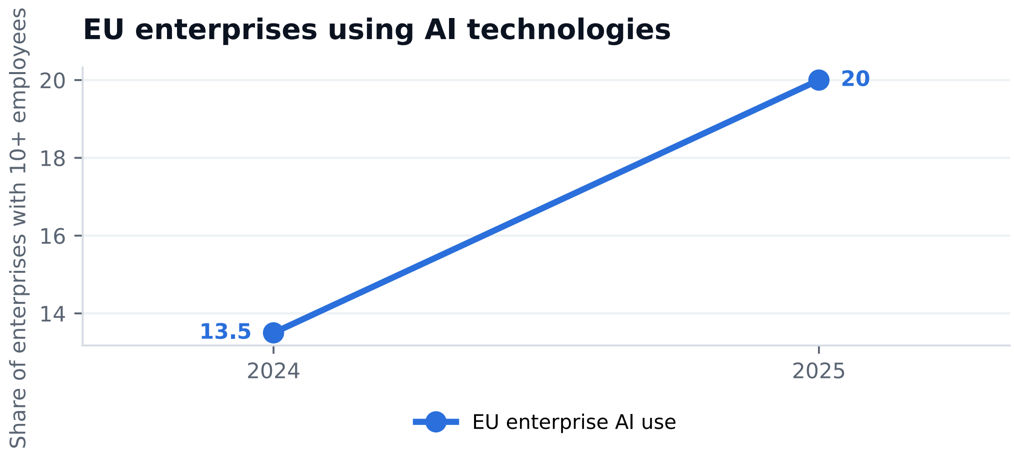

Eurostat, 11 Dec 2025 (official statistic). AIMonger redraw.

Figure 2. Share of EU enterprises (10+ employees) using AI technologies. In short: the field is open; panic closes the window on measured absorption. **Board reading:** your Milan competitor may still be on pilot one, so tempo beats theatre. Source: Eurostat, 11 December 2025 (official statistic). AIMonger redraw.

Stanford's AI Index 2026 compiles similar organisational adoption figures from surveys including McKinsey. Treat the 88 per cent regular-use statistic as McKinsey-sourced, not independent confirmation. Agent deployment in production remains single-digit in compiled figures. The scarce race is production absorption, not chatbot access.

Signals you are in the false race

- **Fifteen pilots, none graduated.** The portfolio grows while no workflow reaches production metrics, proving activity substitutes for absorption in steering reports. A pilot that never graduates is a budget line that never closes, and steering loses the ability to tell progress from motion. The tell is a portfolio deck with go-live dates but no cycle-time or cost-per-case movement.
- **Success defined as launched demos.** Steering celebrates go-live events without cycle-time, quality, or cost-per-case movement finance can tie to the P&L. A demo launch is a moment, not a metric, and celebrating moments trains the firm to produce more of them instead of outcomes. The fix is a board pack that asks for before-and-after numbers on every launch.
- **Strategy decks led by model brand names.** Roadmaps name GPT-5.x, Claude Opus, Gemini, Copilot, DeepSeek-V4, GLM-5.2, or Qwen instead of named workflows, owners, and ninety-day graduate criteria. A strategy that lists model brands, including Chinese open weights as a fashion statement, is a procurement list dressed as a plan. The board should ask for the workflow name, the owner, the placement, and the metric, not the model version.
- **Staff anxiety without training paths.** Employees fear replacement while the firm offers no certification on the governed tools that will actually ship to production. Anxiety without training is a signal that the firm is consuming the narrative without building the capacity to act on it. A training path with dates and completion targets replaces fear with practice.
- **No kill criteria on pilots.** Experiments lack stop dates and failure thresholds, so “monitoring” becomes an indefinite spend line with no decision rule. A pilot without a kill date is a budget commitment in disguise, and the disguise is the word “learning.” Kill criteria turn learning into a decision, not a perpetual line.
- **Shadow AI rising because official AI is theatre.** Staff paste company data into consumer tools when sanctioned programmes are too slow or too shallow to use daily. Shadow AI is the symptom, and official AI being theatre is the cause. When staff route around the sanctioned tool, the sanctioned tool is the problem, not the staff.

Shadow AI means: staff using consumer tools on company data because internal programmes are too slow or too shallow. **On Monday:** contracts pasted into a public chatbot while the sanctioned discovery pilot waits for legal review month four.

Deloitte’s 2026 survey finds only 25 per cent of organisations moved at least 40 per cent of AI experiments to production, while 30 per cent redesigned key processes for AI. Speed theatre without graduation criteria is delay wearing a jersey.

Strategic patience with operating urgency

Patience (what you refuse)	Urgency (what you accelerate)
Fifty parallel pilots	Three named workflows
Model release FOMO	Data-access decisions in weeks
Infinite ”monitoring”	Ninety-day graduate/kill gates
Innovation theatre budget	Operator training this quarter
Vendor-led strategy	Peer-relative tempo metrics

This is not calm for its own sake. It is how high performers behave when the market shouts.

Absorption capacity, put simply: whether the firm can route AI into daily workflows so cost, cycle time, or quality moves, not whether licences were bought. **Operating case:** invoice matching runs overnight with exceptions only; AP no longer chases PDFs by email.

Strategic patience is not caution as a brand pose. It is the discipline to make data-access decisions quickly, train operators now, and refuse a portfolio that exists mainly to reassure the board that something is happening.

Evidence base: the race narrative overstates scarcity of access

Source	Finding	Narrative correction	Does not prove
Eurostat 2025	20.0% EU enterprises (10+) used AI; 13.5% in 2024	Many peers are early	Sector-level leadership
McKinsey 2025	~88% regular use; 39% any EBIT; ~6% >5% EBIT	Access is not advantage	Causal EBIT attribution
Deloitte 2026	25% moved >=40% experiments to production	Pilots without gates stall	Exact Malta sample
Gartner	>40% agentic projects forecast cancelled by 2027	Panic portfolios die	All cancellations are "patience" failures
Stanford Index 2026	Compiled adoption high; agents low	Production scarce	Independent of McKinsey for 88%

Plain read: headlines sell speed; P&L responds to depth. Replace the race with tempo metrics.

Talent and culture

Public narratives often scare staff into believing the profession vanished. Leadership owes a precise story: mechanical cognitive work is repricing; judgement, accountability, and domain craft remain scarce; the firm will train people to operate with AI rather than compete against a chatbot myth.

Workshops matter here. They replace doom-scroll narratives with practice on the actual production system, not a keynote reel.

Example for an HR director in a Valletta professional services firm: run operator certification on the governed discovery tool, and track completion like any compliance training.

Competitive tempo without panic

Useful competitor questions:

- **Which peer workflows actually changed cycle time?** Ask for evidence in delivery or operations, not press releases, because vanity launches rarely include a before-and-after metric. A peer that cut quote cycle time by 30 per cent has a number; a peer that "launched an AI initiative" has a headline. The board should ask for the number.
- **Are they hiring operators and trainers or only buying seats?** Headcount for people who run and override AI signals absorption; licence bulk buys alone often indicate theatre. A firm that hired three operators and certified twelve is building capacity; a firm that bought 400 seats and hired nobody is performing. The org chart is more honest than the press release.
- **Do customer commitments mention AI features that imply production use?** Contractual language about automated turnaround or AI-assisted delivery is a stronger signal than a keynote demo. A SLA that promises four-hour turnaround using AI is a production claim a lawyer reviewed. A demo video is a marketing claim nobody did.
- **Are they leaking shadow-AI incidents in the market?** Rumours of data pasted into consumer tools or audit findings suggest official programmes failed to absorb work staff still need done. A

peer with a shadow-AI incident has an official programme that was too slow or too shallow to use. That is weakness wearing the costume of adoption.

Ignore vanity launches. Copy absorption behaviours that show up in delivery.

Peer tempo scorecard

Compare yourselves to named peers on:

Dimension	Your value	Peer signal	Gap action
Workflows in production with metrics			
Median permission decision days			
Certified operator coverage			
Shadow AI retirement rate			
Cost per successful case			

If you cannot fill the peer column, you are racing headlines, not competitors.

Internal communication pack

Leaders should be able to say, without hype:

1. **We are not chasing every model release.** The firm evaluates models on holdout workflows on a schedule, not on every vendor headline or social-media benchmark thread. This line tells staff that model brand is not strategy, and it protects the team from FOMO-driven context switching. A schedule also means evaluation happens on the firm's timeline, not the vendor's.
2. **We are absorbing N workflows this half.** Name the short list with owners, so staff know where AI effort concentrates instead of guessing from fifteen parallel pilots. A named list with owners turns "AI transformation" from a slogan into a work plan. Staff who know where the effort goes can contribute; staff who guess cannot.
3. **Here are the metrics.** Cycle time, exception rate, or cost per successful case appear in the same pack as the AI narrative, tying spend to outcomes finance recognises. Metrics in the board pack are what separates a programme from a press release. Without them, the AI narrative is a story finance cannot reconcile with the P&L.
4. **Here is the training path.** Operator certification has dates and completion targets, replacing fear-based rumour with practical skill on the production system. A training path with dates is how anxiety becomes capacity. Staff who are trained on the governed tool do not paste company data into a consumer one.
5. **Here is what we will not fund.** Explicit red lines (scattershot pilots, untagged tools, innovation theatre) reduce false-race anxiety more than another keynote reel. A red line tells staff what the firm refuses, which is often more clarifying than what it funds. The list of "not funded" is the discipline that keeps the "funded" list honest.

This reduces false-race anxiety more than another all-hands with a keynote reel.

Portfolio math

If you can deeply absorb three workflows per half-year, a two-year horizon yields roughly twelve operating capacities, more valuable than forty demos.

Patience is a portfolio identity: depth over logo count.

On Monday: half one absorbs quote drafting, policy discovery, and AP exceptions; half two adds onboarding and support draft-and-gate, not twelve chatbots.

Capital allocation under patience

Target portfolio shape for a mid-market firm:

- **Fifty to sixty per cent Absorb.** Fund workflow redesign and production paths that move cycle time, quality, or capacity, the only spend category that can eventually touch EBIT. Absorption spend is what turns a licence into an operating capacity, and it is the only category that pays back. Without it, the rest of the portfolio is overhead.
- **Fifteen to twenty per cent Trust.** Invest in evaluation, provenance, security, and access controls so absorbed workflows survive audit and do not become tomorrow's incident. Trust spend is insurance for the absorption spend, and skipping it is how a working workflow becomes a regulator headline. A workflow that cannot survive audit cannot scale.
- **Fifteen to twenty per cent Learn.** Pay for training, feedback loops, and operator certification so staff run governed tools instead of routing around them in shadow AI. Learn spend is what builds the human capacity to operate the absorbed workflows, and without it the tools sit unused. Operators who are not trained do not absorb.
- **Ten per cent Explore.** Time-box experimental bets with kill dates and pre-registered success criteria, preventing "monitoring the landscape" from consuming the portfolio in disguise. Explore spend is real, but it is capped and time-boxed, not open-ended. An experiment without a kill date is a budget line in costume.

Invert that shape and the false race returns.

Board red lines

- **No pilot without owner and kill date.** Every experiment names a business owner and a calendar stop with graduate-or-kill criteria, preventing indefinite "learning phases" that consume budget without decisions. The owner is who answers for the result, and the kill date is when the experiment stops being an experiment. A pilot with neither is a permanent spend line.
- **No tool purchase without workflow attachment plan.** Licences and platforms attach to named workflows with baselines, so procurement cannot buy seats that float without an absorption target. A licence without a workflow is a seat count, not an operating capacity. The attachment plan is what turns a purchase into a programme.
- **No innovation theatre budget without graduation reporting.** Steering sees which pilots graduated, stalled, or died each quarter, not only which demos launched to applause. Graduation reporting is how the board tells progress from motion, and it is the only honest way to close a pilot. A demo count without a graduation count is theatre.
- **No staff fear without a training response.** If anxiety rises in town halls, operator certification and practical workshops ship on the same timeline as the AI announcement. Fear without training is a leadership failure, and the fix is capacity, not reassurance. The timeline match is what proves the training is real.
- **No board pack with token spend but no cycle time.** Finance reports cost per successful case and workflow tempo alongside licence counts, so the false race cannot hide behind usage metrics alone. Token spend without cycle time is a cost without an outcome, and the board cannot steer on costs alone. The cycle-time column is what makes the spend legible.

Worked example: Madrid logistics mid-market

Composite: 320 staff, family-owned, EU clients, legacy TMS plus Microsoft stack.

False race path: buy Copilot for all, launch five vendor pilots, announce "AI transformation" in press release, with no workflow owner and no metrics. Shadow AI rises in operations within sixty days.

Patience path: 1. **Pick three workflows with business owners.** Customs doc assembly, carrier exception triage, and customer status drafts each get a named owner and baseline cycle-time or rework metric. Three workflows with owners is a programme; fifteen pilots without owners is a panic portfolio. The baseline is what the ninety-day gate measures against. 2. **Assign ninety-day gates on cycle time and rework.** Pre-register what improvement or kill threshold must be met at day ninety before any of the three expands scope or budget. The gate is a decision date, not a review date, and pre-registration stops the team from moving the goalpost. A gate that cannot kill is not a gate. 3. **Fast-track read-only data access with human gates on sends.** Legal-approved read corpora feed discovery and drafting while every external customer communication still passes a staffed human gate. Read-only access is fast because the blast radius is low, and the human gate on sends keeps the blast radius bounded. The combination is how a patient firm moves fast on data and slow on risk. 4. **Train twelve operators and kill one flat pilot at day ninety.** Certification creates operating capacity, and killing the weakest pilot proves patience includes stop discipline, not only start enthusiasm. A programme that can start but cannot kill is not patient, it is indulgent. The kill is the evidence that the gates are real.

Board report: one graduated workflow at 22 per cent cycle-time reduction; two in continue; zero net new licences until graduate one proves EBIT path conversation.

Linking patience to durable advantage

Commodity models punish firms that confuse motion with position. Strategic patience funds reality feedback, retention, and human capacity, while false-race portfolios fund wrappers that repricing will erase.

McKinsey high performers invest in workflow redesign and senior ownership, a patience profile with operating urgency, not a slower clock.

Counter-position: patience is cover for inaction

Another board might argue patience excuses deferral while peers pull ahead. Valid risk. Distinguish patience from delay:

Patience	Delay
Named workflows, owners, dates	"Monitoring the landscape"
Ninety-day graduate/kill	Indefinite pilots
Training scheduled	Awareness events only
Metrics in board pack	Principles slide only

A two-workflow absorption plan with ninety-day gates is patience. "Monitoring the landscape" with no portfolio is delay.

CEO decision criteria

1. **Can you name three workflows absorbing this half with baselines?** Patience requires a short list with owners and before-metrics, not a principles slide about "AI transformation." A named workflow with a baseline is a programme; a principles slide is a poster. The baseline is what the gate measures against.

2. **Does every pilot have a kill date visible to the board?** Steering packs show calendar stops and graduate thresholds, distinguishing patience from open-ended "monitoring the landscape." A kill date visible to the board is a commitment the team cannot quietly drop. A pilot without a visible kill date is a permanent budget line.

3. **Is training completion tracked for production operators?** Certification rates appear beside licence counts, proving the firm is building capacity to run governed tools daily. A licence count without a certification rate is a purchase, not a capacity. The rate is what tells the board whether the tools are being used.

4. **Does finance see cost per successful case, not only licence count?** Unit economics by workflow appear monthly, so OpEx rises only where outcomes justify continuation or expansion. Cost per case is what makes the spend legible to the P&L, and licence count alone is what hides the false race. The unit economic is the metric finance can steer on.

5. **Can you compare tempo to two named peers on one metric?** Peer-relative cycle time or rework beats headline anxiety about hyperscaler spend you cannot replicate or need to replicate. A named peer with a named metric is a comparison; a vague "the market" is anxiety. The peer metric is what anchors the tempo narrative.

If three or more answers are no, you are likely in the false race.

Patience versus delay test

If your plan has dates, owners, and metrics, it is patience. If it has only principles and "monitoring the space," it is delay.

Closing position

You are not behind a universal finish line. You are either absorbing AI into a few workflows with metrics, or you are performing a race for an audience that cannot see your P&L.

Choose tempo. Refuse panic.

When Eurostat shows most EU enterprises still forming AI habits and McKinsey shows EBIT impact concentrated in a small high-performer cohort, the winning move is not to sprint faster. It is to absorb deeper than your peer set on workflows customers feel.

References

1. Eurostat, "20% of EU enterprises use AI technologies," 11 December 2025 (20.0% in 2025; 13.5% in 2024; large enterprises 55.0%). https://ec.europa.eu/eurostat/en/web/products-eurostat-news/w/ddn-20251211-2

2. McKinsey & Company / QuantumBlack, "The State of AI: Global Survey 2025." https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

3. Stanford Institute for Human-Centered Artificial Intelligence, "The 2026 AI Index Report," Economy chapter (organisational adoption figures compiled from surveys including McKinsey). https://hai.stanford.edu/ai-index/2026-ai-index-report/economy

4. Deloitte AI Institute, "The State of AI in the Enterprise: The Untapped Edge" (2026 edition); survey of 3,235 leaders, Aug-Sep 2025. https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html

5. Gartner, "Gartner Says More Than 80% of Enterprises Will Have Used Generative AI APIs or Deployed Generative AI-Enabled Applications by 2026," press release, 11 October 2023. https://www.gartner.com/en/newsroom/press-releases/2023-10-11-gartner-says-more-than-80-percent-of-enterprises-will-have-used-generative-ai-apis-or-deployed-generative-ai-enabled-applications-by-2026

6. Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," press release, 25 June 2025. https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

7. Regulation (EU) 2024/1689 of the European Parliament and of the Council (EU AI Act). https://eur-lex.europa.eu/eli/reg/2024/1689/oj

8. OECD.AI policy materials on enterprise AI adoption and governance tempo. https://oecd.ai/en/ai-principles

Published by AIMonger . https://aimonger.com . Malta . Europe