AIMONGER
WHITE PAPER

AI Discovery

Absorption Capacity: Why Access to AI Is Not Advantage

The AI question is no longer whether the tools exist. Eurostat reports 20.0 per cent of EU enterprises using AI in 2025, while global surveys show access rising much faster than EBIT impact. The board variable is absorption capacity: whether cheaper intelligence can be routed into workflows that change cost, cycle time, and quality.

David Saliba . 2026-07-02 . aimonger.com/whitepapers/absorption-capacity-ai-advantage/

The problem in one sentence

Every serious competitor can buy comparable model capability. The returns go to companies that absorb that capability into how they operate.

Walk into a boardroom in 2026 and AI is on the agenda. Walk into operations of the same company and you often find something stranger: a few Copilot seats, a finance pilot, a local champion in operations, and a cost structure that still looks like 2022. AI is present on paper. It has changed little that a CFO would recognise.

What you see globally right now is a strange split. US hyperscalers are spending at infrastructure scale, Silicon Valley still celebrates seat-count adoption, and enterprise vendors are adding AI buttons everywhere. Gartner projected that by 2026 more than 80 per cent of enterprises would have used generative AI APIs or deployed GenAI-enabled applications in production environments, up from less than 5 per cent in 2023. McKinsey's State of AI 2025 survey found that nearly nine in ten respondents say their organisations regularly use AI in at least one function. Those figures describe access and experimentation. They do not describe advantage.

McKinsey also reports the uncomfortable companion number: only 39 per cent of organisations attribute any enterprise-level EBIT impact to AI, and most of those attribute less than 5 per cent of EBIT. About 6 per cent qualify as high performers, organisations that attribute more than 5 per cent of EBIT to AI and report significant value. The pattern that separates them is not access to a better model. It is redesign of workflows, faster scaling of what works, and senior ownership of outcomes.

This paper is about that gap as it appears from a CIO office: model access is global, but absorption is local. It depends on the systems, permissions, managers, and operating cadence inside the firm.

> **Absorption capacity, plainly:** whether the company can put widely available model capability into daily workflows so cost, cycle time, or quality actually moves, not whether licences were bought. **In practice:** a German logistics group routes supplier invoices through extraction overnight; accounts payable clears exceptions only, instead of three staff chasing PDFs in email.

> **High performers (>5% EBIT), in short:** McKinsey's label for the roughly 6% of surveyed organisations that attribute more than 5% of enterprise EBIT to AI and report significant value, not firms that merely bought more seats. **Worked case:** if Copilot appears in OpEx but consolidated EBIT margin is flat, you are not in this cohort yet; the gap is workflow redesign, not model brand.

This is no longer a chatbot story

The chat window was the first consumer wrapper around large models, and it still shapes how many executives picture the technology: a clever assistant you type at. That picture is now incomplete. The interface was never the point. The point is that reading, writing, classifying, extracting, summarising, and drafting, the mechanical middle of knowledge work, can be performed by software at a collapsing marginal cost.

Consequential deployments in 2026 rarely look like open-ended chat. They look like:

- An accounts-payable pipeline that extracts line items from tens of thousands of supplier invoices, flags the small share that disagree with purchase orders, and routes only those to a human. The work that used to absorb three full-time staff now runs overnight, and the humans spend their day on the cases where judgement actually changes the outcome.
- An internal discovery system that answers "have we dealt with this clause before?" from decades of contracts in seconds instead of a multi-day email chase. A specialist can confirm or refute a position before the client meeting ends, rather than promising a reply by Friday and losing the window.
- A support operation where the model drafts a resolution, checks it against policy, and escalates ambiguous cases with a summary attached. Agents no longer type boilerplate from scratch; they review, correct, and send, which lifts volume per agent without lowering quality.
- An engineering practice where maintenance work, dependency upgrades, test scaffolding, and migration boilerplate are generated and reviewed rather than typed from scratch. Senior engineers spend their hours on architecture and review, and the repetitive scaffolding that used to block feature work gets cleared in hours.

None of this is science fiction, and none of it is "a chatbot project." It is workflow plumbing. It becomes decisive in the same way electricity became decisive: invisible, load-bearing, and expensive to be late on.

Eurostat's enterprise AI figure adds a second layer. In 2025, 20.0 per cent of EU enterprises with 10 or more employees reported AI use, up from 13.5 per cent in 2024. That is still a minority of firms, but among firms that compete internationally, the relevant peer set already has access. Waiting for "AI to arrive" is the wrong clock. The clock is how fast peers absorb what is already available.

The model is the engine. The business system is the vehicle.

> **Workflow redesign means:** changing who does what, in which system, with which data, not adding an AI button to an unchanged process. **On Monday:** a Malta professional-services firm stops asking associates to rebuild contract summaries from scratch; the system drafts from the archive and the partner reviews only material deviations before client send.

Frontier models matter. An engine on a workbench moves nothing. What moves is a vehicle: engine plus transmission, steering, brakes, fuel, and a driver who knows the route.

In enterprise terms:

Vehicle part	Business meaning
Transmission	Integration with systems of record (ERP, CRM, document stores, case systems)
Steering	Domain-specific workflow design: where the model sits in the process
Brakes	Evaluation, review, escalation, and human gates on consequential decisions
Fuel	Data that is accessible, permissioned, and current enough to be useful
Driver	A named owner accountable for outcomes, not only for procurement

This distinction explains a puzzling fact: model capability has improved for years while many corporate AI returns have not. The engine is rapidly becoming a commodity. Several vendors will sell you a very good one, and as of July 2026 competitive open weights (DeepSeek-V4, Zhipu GLM-5.2, Qwen3.x) sit beside closed GPT-5.x / Claude Opus-class / Gemini 3.x APIs. The price per unit of capability keeps falling, including on private trays. The vehicle, the business system wrapped around the engine, cannot be purchased as a single SKU. It has to be built from your data, your processes, your permissions, and your people.

Meanwhile the global narrative is still mostly about bigger models, bigger data centres, and which national lab shipped the latest open weight. In a mid-market firm, the decisive question is usually smaller and more practical: can the system read the right archive, change the right workflow, and get approved by the right owner this quarter? Chinese open models change the procurement shortlist and the sovereignty options. They do not replace absorption.

That is where durable advantage now lives, precisely because competitors cannot copy it by signing the same contract.

The new divide: who can absorb AI

For twenty years the strategic question about new technology was “who has it?” For generative AI that question is nearly worthless, because the answer for serious enterprises is converging on “everyone who wants it.”

The question that separates companies is: who can absorb it?

> **Operating capacity, put simply:** a production workflow with monitoring, escalation, owners, and a metric finance recognises, not a sandbox demo that survives only while the project team is in the room. **Operating case:** claims triage runs every morning against the case system; exceptions queue to a named adjuster with provenance attached; rollback is one release tag away.

Absorption capacity is concrete. It means:

1. **Data access.** Someone can grant a system permission to read the contract archive or case store without a six-month security theatre loop that never ends in a decision. Data access is the first gate because every downstream capability, retrieval, grounding, drafting, depends on the model seeing the right records. A mid-market insurer we worked with in principle spent eleven months waiting on a single archive permission; the pilot expired before the data arrived.

2. **Process clarity.** Workflows are documented well enough that you can point at the step where a model should sit and defend that placement in a steering meeting. Without a written process, the team argues about where AI belongs instead of wiring it in, and the pilot drifts into a demo that never touches production.

3. **Manager posture.** Middle managers treat AI-assisted workflows as capacity to redeploy, not as a personal threat to be quietly starved of data and headcount. When a manager fears the tool will cost the team’s seats, the pilot receives thin data and no airtime, then fails for reasons that look like the model’s fault.

4. **Operating review tempo.** Legal, security, and the business can decide on a bounded use case in weeks with clear risk boundaries, rather than deferring indefinitely to the next quarter. Tempo

here is a governance choice: a fast path for low-risk cases, and a slower, documented path for consequential ones.

5. **Outcome ownership.** A senior person owns the result and can change the process, not only buy another tool or extend a pilot licence. Ownership means a name on the metric and the authority to alter staffing, prompts, and escalation rules when the numbers demand it.

Two companies with identical AI budgets will diverge on these dimensions. One deploys document intelligence into claims or procurement in a quarter because the workflow was mapped, data was governed, and operations wanted it. The other buys the same product and eighteen months later owns a demo environment nobody logs into, because every integration request died in a queue.

Absorption capacity deserves the same board scrutiny as balance-sheet leverage. It determines how fast the company can convert a falling input cost into margin or growth.

The collapsing cost of cognitive work

> **Cognitive work (plain English):** the mechanical middle of knowledge jobs, reading, extracting, drafting, classifying, summarising, as distinct from judgement, accountability, and relationships. **On the P&L:** if first-pass invoice coding cost drops tenfold, the finance question is whether you review all lines or still sample 5% because that was the old constraint.

Underneath the tooling conversation is a repricing. The unit cost of a large class of cognitive tasks, first-draft research, document review, data extraction, report assembly, routine analysis, standard correspondence, boilerplate code, is collapsing. Not to zero, and not for every task. Judgement, accountability, relationships, and taste remain expensive. But the mechanical middle of knowledge work is being repriced the way spreadsheets repriced ledger arithmetic.

Be careful what this does and does not imply for labour. It does not mean analysts, lawyers, or engineers disappear. Spreadsheets did not eliminate accountants; they eliminated arithmetic and multiplied what one accountant could oversee. It does mean that a workflow designed around expensive human reading, a team sampling 5 per cent of transactions because reviewing all of them was unaffordable, is designed around a false constraint.

When review becomes cheap, you review everything and staff humans on the exceptions. When drafting becomes cheap, the bottleneck moves to deciding and verifying. Companies that redesign work around the new cost structure gain tempo. Companies that keep the old staffing model and add a chatbot on the side pay for both.

The nearer threat is tempo, not replacement

Public debate often frames AI as a replacement story. For most mid-market firms, the nearer competitive threat is quieter: being out-executed by a peer with similar headcount and better workflows.

Across the US and Asia, the public story often sounds like a labour-market shock or a platform race. In many mid-market sectors, the boardroom issue is more prosaic: the competitor that quotes faster, clears exceptions earlier, and keeps senior people out of chase work.

Tempo shows up as:

- **Quote-to-cash tempo.** Shorter cycles emerge when document assembly and chase work collapse, so sales teams spend less time waiting on paperwork before revenue lands. A tender response that used to take five days of drafting can leave the door in two, which is often the difference between being first or second into the buyer's shortlist.
- **Product onboarding speed.** New products reach the field faster when policy and procedure discovery is no longer tribal knowledge trapped in senior inboxes. A new hire can find the

governing procedure in minutes instead of scheduling three coffees with people who might remember it.

- **Exception quality.** Humans see only the cases that need judgement, which raises the quality of decisions on the exceptions that actually move margin or risk. When the model clears the routine, reviewers stop fatiguing on easy cases and bring full attention to the marginal ones.
- **Cost-to-serve without headline cuts.** Capacity is absorbed into growth or service levels rather than a visible headcount reduction that frightens the board before metrics prove out. This avoids the political risk of a redundancy programme that precedes the evidence that the tool works.

Replacement narratives generate fear and freeze decisions. Tempo narratives force operating choices: which workflows get redesigned this quarter, who owns them, and what metric proves the redesign worked.

McKinsey's high-performer cohort, roughly 6 per cent of surveyed organisations, reports pushing for transformative innovation via AI, redesigning workflows, scaling faster, and investing more. That is a tempo description, not a model-shopping description.

Intelligence is not automation. Know which one you need.

Boards often confuse two different buys:

Mode	What you buy	What success looks like
Automation	Deterministic rules, RPA, fixed pipelines	Same output every time; audit trail of rules
Intelligence	Models that draft, classify, retrieve, propose	Better average quality and coverage; human gates on consequential steps

Automation fails when the world is messy. Intelligence fails when nobody owns the mess, when there is no evaluation set, no escalation path, and no person who can change the prompt, the retrieval corpus, or the process.

Many "AI projects" are automation projects wearing a model badge. Many "automation projects" would benefit from a bounded intelligence layer for the exception path. Mixing the two without naming the difference produces dashboards that look busy and P&L lines that do not move.

The wrong question, and the right one

Wrong question: Which AI tool should we buy?

Right question: Which parts of the business should become more intelligent this year, and who owns the operating change?

Tool selection is a procurement decision. Absorption is an operating decision. Start with workflows that burn time on chase work, document assembly, cross-system lookups, and repeated approvals. Ask whether the data exists, whether permissions can be granted, whether a human gate is defined for high-risk actions, and whether a metric can show progress in ninety days.

If those answers are missing, buying another licence does not create capacity. It creates inventory.

From experiments to operating capacity

Useful AI programmes pass through stages that boards can recognise:

1. **Sandbox.** People learn what models can do while nothing production-critical depends on them, so experimentation does not put customer outcomes at risk. The sandbox is for building intuition about failure modes, not for claiming adoption in the board pack.

2. **Bounded pilot.** One workflow runs with a named owner, evaluation criteria, and a rollback plan so the board can see whether absorption is real before scale funding. The pilot's job is to produce a metric, not a story.
3. **Operating capacity.** The workflow runs in production with monitoring, escalation, and a cost or quality metric on a management dashboard that finance recognises each month. At this stage the workflow keeps running when the champion takes holiday.
4. **Compounding.** Adjacent workflows reuse the same data permissions, evaluation habits, and integration patterns so each graduate lowers the cost of the next absorption move. A governance pattern that worked for claims becomes scaffolding for procurement.

Most stuck programmes never leave stage 1 or 2. They collect demos. High performers industrialise stage 3 and treat stage 4 as a portfolio, not a hope.

Gartner's projection that more than 80 per cent of enterprises would use GenAI APIs or GenAI-enabled applications by 2026 describes stages 1 to 3 becoming common. McKinsey's EBIT figures show that stage 3 at enterprise scale remains uncommon. The board's job is to refuse a permanent sandbox culture.

Leadership changes shape, not owner

Absorption does not require the CEO to become a machine-learning engineer. It requires the CEO and board to insist on:

- **Priority workflow short list.** The board funds a short list of priority workflows rather than fifty pilots, so capital concentrates where chase work and document load are measurable. Spreading budget across a long tail of pilots guarantees none of them gets the integration and owner attention needed to graduate.
- **Named owners with authority.** Each workflow has an owner with authority to change process, not merely to report status upward when integration tickets stall. Status reporting without authority is how a pilot stays green on a slide while the queue behind it never clears.
- **Time-boxed legal and security review.** Legal and security review bounded cases in weeks with explicit risk boundaries, so low-risk absorption moves do not die in indefinite queue. A bounded fast path is not a lower standard; it is a published one, with a slower track reserved for genuinely consequential cases.
- **Finance-recognised metrics.** Metrics that finance recognises appear on the management dashboard, so AI spend connects to cycle time, exceptions, or hours removed rather than seat counts alone. When finance cannot read the metric, the board is voting on faith.
- **Operator training, not demo theatre.** Training treats staff as operators of new capacity with practice on the real system, not as spectators who applaud a town-hall demo and revert to email. A team that has only seen the demo will revert the moment the project team leaves the room.

Workshops and team training matter here as operating infrastructure, not as soft extras. Tools without trained operators become shelfware. Operators without redesigned workflows become frustrated early adopters who revert to email.

What boards should do in the next ninety days

1. **Inventory reality.** List AI licences, pilots, and production workflows, then separate chat access from workflow absorption so the board sees what actually changed how work runs. The inventory is the first defence against licence-counting as adoption, because it forces a honest count of workflows that moved a metric.

2. **Pick three workflows.** Choose processes with measurable chase work or document load, and reject "strategy AI" until at least one operational win exists with a baseline metric. Three is deliberate: it concentrates ownership and review attention without sprawling into a portfolio no one steers.
3. **Assign owners.** Each workflow needs a business owner and a technical counterpart under one accountable name, because dual ownership without a single signer is how pilots stall for quarters. The accountable name signs the graduate or kill decision.
4. **Set absorption gates.** Write the data-access decision, evaluation criteria, human gate definition, and ninety-day metric before more spend, so procurement cannot outrun operating readiness. A gate written before spend is a contract; written after, it is a justification.
5. **Kill or graduate.** Every pilot either graduates to operating capacity with a dashboard metric or is closed with a written reason, because indefinite pilots are a tax on attention and budget. A closed pilot with a documented reason is evidence of discipline, not failure.
6. **Train the operators.** Schedule structured enablement for the teams who will live in the new workflow, and do not confuse a town hall with operator practice on production paths. Operators who have drilled failure modes will keep the workflow alive when the champion rotates.
7. **Report tempo.** The board pack shows cycle time, exception rate, and hours removed rather than token counts or model brand names, so directors judge absorption not experimentation theatre. Tempo metrics let the board compare itself to peers on what matters: speed of operating change.

Cheaper intelligence, faster learning, better memory

> **AI discovery, briefly:** finding the right policy, clause, prior case, or operating fact at the moment of work, institutional memory with a searchable interface and audit trail. **Operating case:** a cross-border insurer answers "have we approved this clause in France before?" from governed archives in seconds instead of a three-day email chase across offices.

Three capacities compound when absorption works:

- **Cheaper intelligence.** Mechanical cognitive work costs less per unit, so coverage expands from sampling a fraction of cases to reviewing everything and staffing humans on exceptions. The firm moves from a posture of "we cannot afford to check every invoice" to "we check every invoice and intervene only where the model disagrees with the record."
- **Faster learning.** Evaluation loops and production feedback improve prompts, retrieval, and process design faster than quarterly vendor releases can substitute for operating knowledge. Each wrong answer that becomes a fixture makes the next quarter's releases safer to adopt.
- **Better memory.** Institutional knowledge becomes discoverable instead of trapped in inboxes and personal drives, which shortens quote cycles and reduces inconsistent answers across teams. A new joiner reaches the same answer a ten-year veteran would have given, without the three-month induction.

AI discovery, finding the right clause, prior decision, policy, or case pattern at the moment of work, is often the highest-value absorption move for professional services, regulated operations, and multi-entity groups. It is not a science project. It is operating memory with a modern interface and an audit trail.

Worked patterns: what absorption looks like in practice

The following patterns are illustrative operating designs, not client case studies. They show where boards should look for measurable change.

Pattern A: Exception-first document operations

Before: A team samples a fraction of inbound documents because full review is unaffordable. Exceptions are found late. Cycle time is dominated by chase email.

After absorption: Every document is extracted and compared to the system of record. Only disagreements route to humans. The metric is exception rate, mean time to clear an exception, and hours removed from sampling work.

Absorption requirements: API or export access to the system of record; a written definition of "disagreement"; a human owner for the exception queue; an evaluation set of known hard cases.

Pattern B: Institutional memory for professional work

Before: "Have we seen this before?" means asking three colleagues and hoping someone remembers a folder name.

After absorption: A governed discovery system returns prior agreements, policies, or case patterns with provenance. Humans still decide. The system shortens search time and raises coverage of institutional memory.

Absorption requirements: Corpus ownership; temporal marking of superseded documents; permission model; holdout evaluation (see the companion knowledge-systems paper).

Pattern C: Draft-and-gate correspondence

Before: Specialists draft every customer or regulator response from a blank page.

After absorption: The system drafts from policy and prior correspondence; a named role edits and sends. The metric is turnaround time and rework rate, not "number of AI emails."

Absorption requirements: Clear send authority; refusal rules when evidence is thin; logging of what was drafted versus what was sent.

Anti-patterns boards should recognise early

1. **Licence counting as adoption.** Seat numbers rise while workflow metrics are unchanged, which lets leadership hear "we have AI" while operations still run like 2022. The fix is to report workflow attachment, not seat count, in the board pack.
2. **Pilot museum.** Twelve demos accumulate with zero graduates and no kill dates, so the firm pays for experimentation inventory instead of operating capacity. Each immortal pilot also consumes steering attention that a graduate would deserve.
3. **Shadow IT intelligence.** Teams paste sensitive documents into consumer tools because internal absorption is too slow, which creates risk exposure the official dashboard never counts. The shadow path grows precisely because the approved path cannot deliver a working workflow this quarter.

> **Shadow IT (GenAI):** staff using consumer chat tools on company data because the approved internal path is too slow, too weak, or too blocked to be useful. **On Monday:** legal staff paste contract excerpts into a public assistant because the internal discovery pilot never received archive permissions; the board sees zero official AI usage while risk exposure rises.

4. **Model bake-off without a workflow.** Months of vendor comparison proceed with no owner for the operating change, so procurement optimises a logo while process design stays untouched. The bake-off produces a winner and no workflow metric.
5. **Automation cosplay.** A brittle script wrapped in model marketing language runs without an evaluation set, which produces confident wrong outputs on the exceptions that matter most. When the script fails, it fails on the case the firm most needed to get right.
6. **Training as a town hall.** Awareness events substitute for operator practice on the real system, so staff revert to email the week after the launch applause fades. A town hall teaches no one to handle a refusal or read provenance under pressure.
7. **Metric substitution.** Token spend or chat sessions get reported instead of cycle time and quality, which lets the board pack look active while P&L lines stay flat. Token counts measure consumption; they never measure absorption.

Each anti-pattern is a capital allocation smell. Name it in steering meetings.

CFO pack: numbers that survive scrutiny

Ask for a one-page insert in the board pack:

Metric	Definition	Cadence
Workflow cycle time	Start-to-finish for the target process	Weekly
Exception rate	Share of cases needing human judgement	Weekly
Hours removed	Estimated hours no longer spent on chase/ sampling/drafting	Monthly
Cost per successful case	Model + retrieval + review minutes, monetised	Monthly
Promotion rate	Pilots graduated vs killed vs still open	Monthly
Training completion	Operators certified on the production workflow	Per release

If the insert cannot be filled, the programme is not ready for scale funding.

Malta and mid-market notes

mid-market firms rarely lose because they cannot buy a model. They lose tempo when:

- **Fragmented data ownership.** Data sits in fragmented systems with unclear ownership, which means every absorption pilot waits on a fresh integration debate instead of reusing governed access paths. The first pilot in a domain pays the integration tax; without reusable paths, every later pilot pays it again.
- **No bounded legal fast path.** Legal review has no bounded fast path for low-risk pilots, so teams choose shadow tools or defer rather than ship a measured workflow this quarter. A published fast path lets low-risk cases move while high-risk cases get the review they need.
- **Undocumented key processes.** Key processes live in people's heads rather than documented workflows, which makes it impossible to point at the step where a model should sit and defend that choice. Undocumented processes also block evaluation, because there is no baseline to measure against.
- **Conference spend over operator enablement.** Training budgets fund conferences instead of operator enablement on production systems, so licences arrive before anyone can run the workflow without the project team. A conference badge does not teach an operator to handle a refusal at 4pm on a Friday.

Malta-origin firms serving EU clients have an additional constraint: reputation travels quickly in a dense professional network. A visible operational win becomes a trust asset. A public failure becomes a trust tax. That asymmetry argues for narrow, well-measured absorption over broad, noisy experimentation.

Eurostat's 20.0 per cent enterprise AI figure for the EU (2025) means many local peers are still early. The window to open a tempo gap is real. It closes as adoption normalises inside your sector.

Operating cadence for the executive team

Suggested monthly rhythm:

1. **Portfolio review.** Review three priority workflows only with status graduate, continue, or kill, so steering meetings stay decision-focused instead of becoming AI theatre. Three workflows fit on one screen and force a decision each month rather than a status read.
2. **Absorption blockers.** Escalate data access, legal, integration, and owner vacancy with named dates, because blockers without deadlines become the polite reason pilots never graduate. A blocker without a date is a deferral in disguise.
3. **Metric movement.** Compare CFO pack metrics to baseline each month, so the executive team sees whether absorption moved cycle time, exceptions, or hours removed. Movement against baseline is the only honest test of whether the spend is working.
4. **Risk and audit.** Review incidents, near-misses, and evaluation failures explicitly, so trust controls improve with production use rather than after a visible customer mistake. A near-miss reviewed this month prevents the incident next quarter.
5. **Training debt.** Track who can operate the system without the project team in the room, because training debt is the hidden reason workflows collapse when champions rotate. An operator who cannot run the workflow alone is a single point of failure.

Quarterly, the board sees the same story in compressed form. No separate "AI theatre" slide that cannot connect to operations.

Appendix: ninety-day absorption charter (template)

Copy into the programme brief before spend increases. Leave the Entry column blank until owners fill it.

Field	Entry
Workflow name	
Business owner	
Technical owner	
Baseline metric (date and value)	
Target metric (90-day)	
Data sources required	
Permission decision date	
Evaluation set size and owner	
Human gate definition	
Training plan (roles and dates)	
Budget ceiling (tools, tokens, and people days)	
Graduate criteria	
Kill criteria	
Board reporting line	

Plain read: charters without kill criteria are wishes. Charters without owners are orphan projects. Empty cells are intentional until filled.

Appendix: questions for the next CIO and CFO meeting

1. Which three workflows are in the absorption portfolio, and which were rejected?
2. What is the oldest open pilot, and why is it still open?
3. Where is shadow IT GenAI occurring, and what internal path would retire it?
4. Which integration or data-access ticket is the critical path?
5. What did we spend last quarter that produced no workflow metric movement?
6. Which operators are trained to run production without the project team?
7. What residual risks have been accepted in writing?
8. Which documentation do we already need because these workflows are now operating dependencies?

Evidence base: adoption is scaling faster than redesign

The adoption-value gap is visible across independent datasets, although the populations differ and should not be blended into a single universal rate.

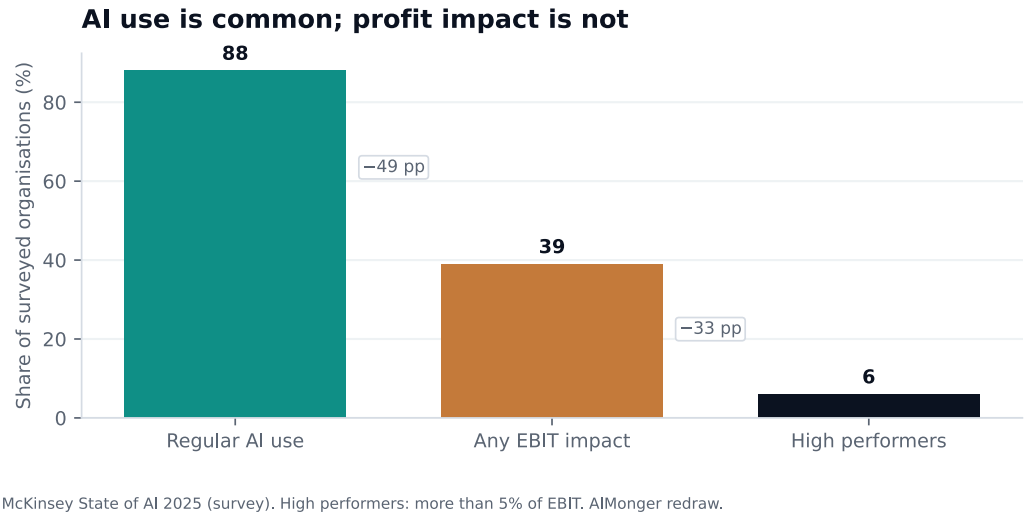

Figure 1. Organisations regularly using AI vs those reporting enterprise EBIT impact. Plain read: most firms have AI somewhere; few can point to profit. **Finance reading:** Copilot seats show up in OpEx; the P&L still looks like last year. Source: McKinsey State of AI 2025 (survey). AIMonger redraw.

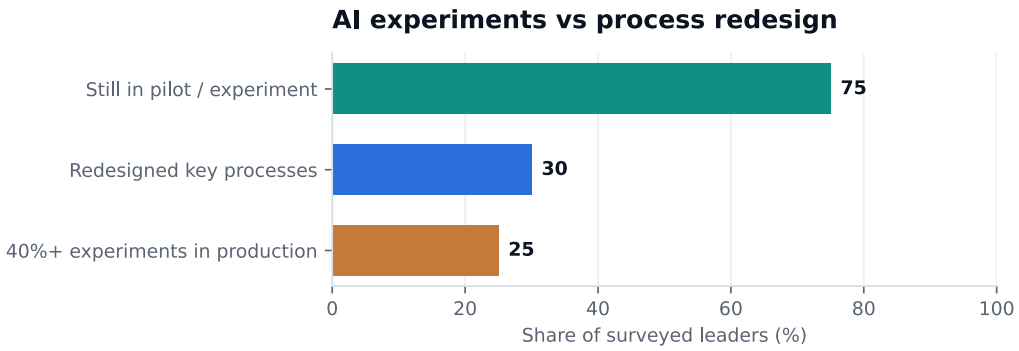

Deloitte State of AI in the Enterprise 2026 (survey, 3,235 leaders). AIMonger redraw.

Figure 2. Share of leaders reporting production scale vs process redesign around AI. In short: many experiments graduate; fewer workflows are rebuilt around AI. **Board reading:** a Benelux manufacturer runs five GenAI pilots but only one changed how orders are released to the floor. Source: Deloitte State of AI in the Enterprise 2026 (survey, 3,235 leaders). AIMonger redraw.

Evidence	Finding	What it does and does not prove
Eurostat enterprise survey (2025)	20.0% of EU enterprises with 10+ workers used at least one AI technology, up 6.5 percentage points from 2024; 17.0% of small, 30.4% of medium, and 55.0% of large enterprises used AI	Representative EU business adoption; does not measure ROI or depth of use
McKinsey State of AI (2025)	88% of surveyed organisations reported regular AI use in at least one function; 39% reported any enterprise EBIT impact; about 6% met McKinsey's high-performer definition	Global executive survey; self-reported impact and a different sample from Eurostat
Deloitte State of AI in the Enterprise (2026)	25% of 3,235 surveyed leaders said their organisation had moved at least 40% of AI experiments into production; only 30% reported redesigning key processes around AI	Directly supports the pilot-to-process gap; respondents were leaders involved in AI initiatives, so the sample is not all firms
Stanford AI Index (2026)	Agent use remained single-digit across nearly all functions; industry produced roughly 90% of notable frontier models in 2025	Cross-source synthesis, useful for technology-market direction; organisational adoption figures in the Index partly compile McKinsey survey evidence

The conclusion is not that every company is failing. It is that access, experimentation, production, process redesign, and financial impact are separate stages. A board dashboard that collapses them into one adoption percentage hides the operating problem.

Methodology note

This paper triangulates four source types: official EU statistics, global executive surveys, a multi-source academic index, and public analyst forecasts. Forecasts describe expected market direction, not observed outcomes. Survey percentages are not directly comparable because samples, sectors, question wording, and definitions of "use" differ. All financial impact claims remain self-reported unless a source states otherwise.

Competing interpretation: AI value may simply need more time

A reasonable counterargument is that enterprise technology has always shown a productivity lag: organisations incur implementation cost before complementary process and skill investments produce measurable surplus. The data supports that possibility. It does not support passive waiting. Deloitte found only 30% redesigning key processes around AI, while PwC's 2025 jobs analysis found productivity

growth strongest in AI-exposed industries. The management implication is to invest in complements, workflow redesign, skills, data access, evaluation, and measure the lag explicitly.

Board decision standard

An absorption programme clears investment committee only when it identifies:

1. **Verified baseline workflow.** The programme names a workflow with a verified baseline dated and signed, so the board can see movement against reality rather than against a slide assumption. A baseline signed by the business owner is the anchor that stops later numbers from drifting into optimism.
2. **Process owner with authority.** A process owner with authority to redesign the workflow can change steps and staffing, not merely sponsor another tool purchase. Sponsorship without authority leaves the process untouched while the licence renews.
3. **Data-permission path with date.** A data-permission path exists with a decision date, so integration and legal queues cannot silently become the reason the pilot never reaches production. A dated permission path converts a vague "we are working on access" into a commitment the board can track.
4. **Leading and financial indicators.** A ninety-day leading indicator and a twelve-month financial hypothesis are written before scale spend, so investment committee can judge tempo and EBIT path separately. The leading indicator shows whether the workflow moves this quarter; the financial hypothesis shows whether it is worth scaling.
5. **Training and review budget.** A training and review-capacity budget is funded alongside tools, because operators and human gates are part of the cost of trustworthy absorption. Funding tools without funding the people who run them is how pilots produce demos and no graduates.
6. **Graduate, pause, or kill criterion.** A graduate, pause, or kill criterion is pre-written with owner and date, so the programme cannot become an immortal pilot that consumes attention without proof. A pre-written kill date is the single cheapest discipline a programme can adopt.

This standard is intentionally stricter than approving a licence. It is also cheaper than funding a portfolio of pilots that cannot explain how value will appear.

Research addendum: infrastructure access is not absorption

OECD economic work on AI diffusion treats data-centre access, skills, innovation capacity, and connectivity as joint drivers of projected gains. Domestic or purchased compute does not produce surplus without evaluation, workflow redesign, cost allocation, and trained operators.

Separately, Gartner's public January 2026 AI spending forecast, roughly **$2.5 trillion** worldwide AI spend in 2026 with infrastructure the largest component, is a market forecast, not an enterprise budget template. It does underline a practical board risk: cheaper tokens and larger infrastructure markets can increase total spend unless workflows and routing discipline convert usage into accepted outcomes.

Closing position

The 2026 divide is not who has AI. It is who can absorb it.

Gartner and McKinsey describe a market where access is mainstream and enterprise-level financial impact is still concentrated. Eurostat shows EU enterprise adoption remains uneven, 20.0 per cent in 2025, up from 13.5 per cent in 2024, which means early absorbers in a sector can still open a tempo gap before peers catch up.

Boards that keep shopping for engines while neglecting the vehicle will fund a permanent demo culture. Boards that treat absorption capacity as a managed asset, workflows, data permissions, evaluation, owners, training, will turn commodity intelligence into operating advantage.

References

1. 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
2. McKinsey & Company / QuantumBlack, "The State of AI: Global Survey 2025." https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
3. 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
4. Stanford Institute for Human-Centered Artificial Intelligence, "The 2026 AI Index Report," Economy chapter. https://hai.stanford.edu/ai-index/2026-ai-index-report/economy
5. PwC, "The Fearless Future: 2025 Global AI Jobs Barometer," 3 June 2025. https://www.pwc.com/gx/en/news-room/press-releases/2025/ai-linked-to-a-fourfold-increase-in-productivity-growth.html
6. Eurostat, "20% of EU enterprises use AI technologies," 11 December 2025 (20.0% in 2025; 13.5% in 2024). https://ec.europa.eu/eurostat/en/web/products-eurostat-news/w/ddn-20251211-2
7. Eurostat, "Use of artificial intelligence in enterprises," Statistics Explained. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Use_of_artificial_intelligence_in_enterprises
8. 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
9. OECD, "AI Meets Trade: Global Linkages and the Cross-Country Distribution of the Gains from AI," March 2026. https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/03/ai-meets-trade_6001acf4/13081644-en.pdf
10. Gartner, "Gartner Says Worldwide AI Spending Will Total $2.5 Trillion in 2026," press release, 15 January 2026. https://www.gartner.com/en/newsroom/press-releases/2026-1-15-gartner-says-worldwide-ai-spending-will-total-2-point-5-trillion-dollars-in-2026

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