HTech / National Enterprise AI Program™

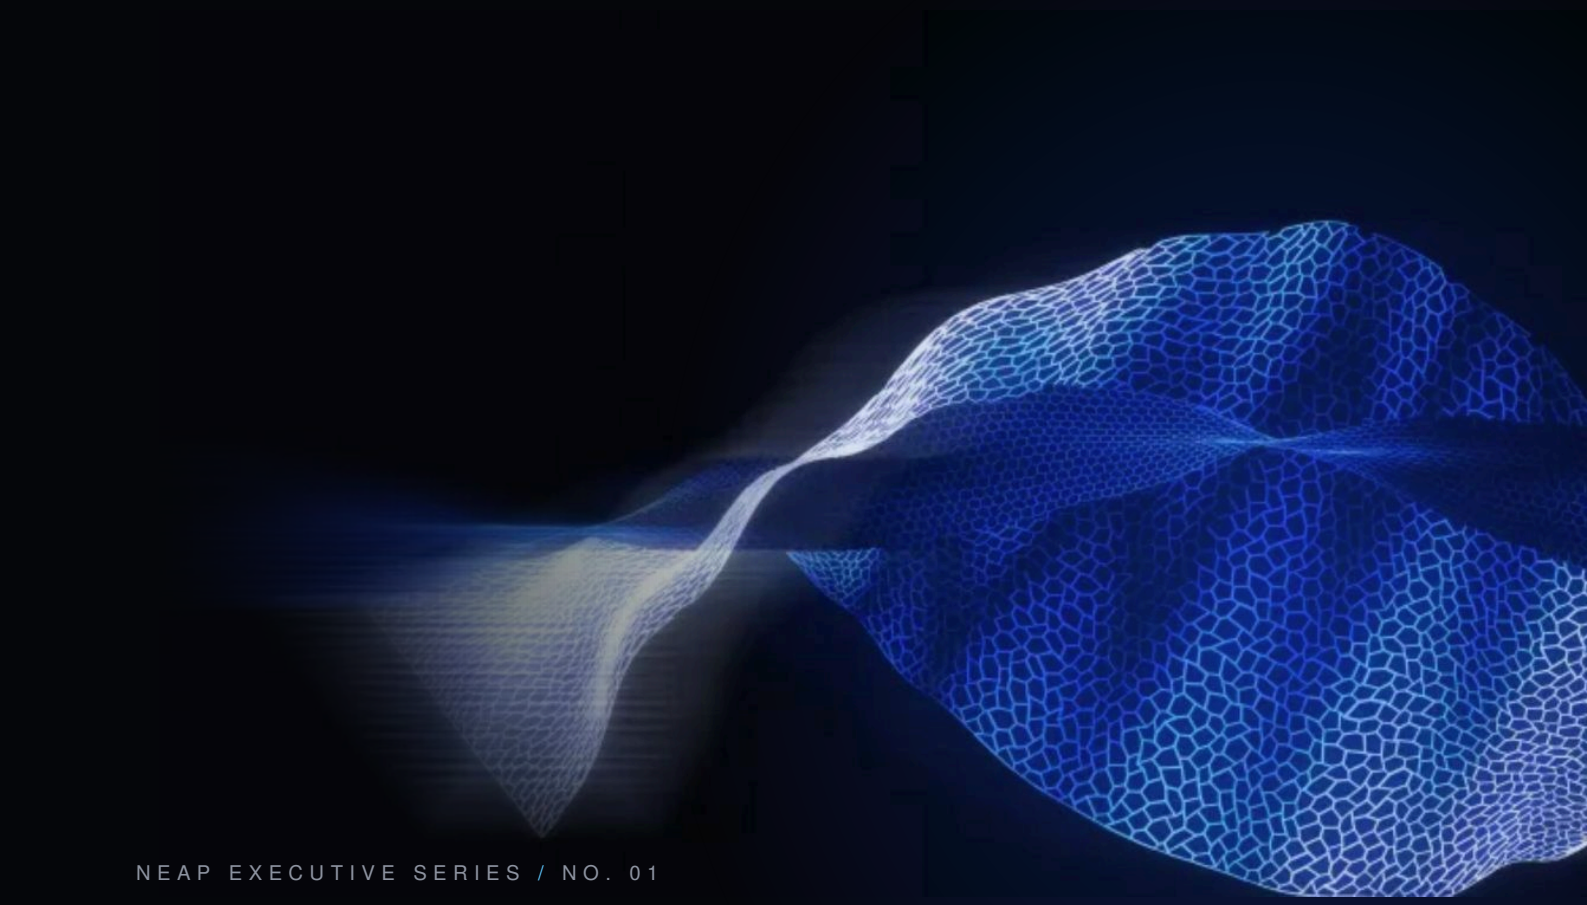

NEAP EXECUTIVE SERIES / NO. 01

Enterprise AI: the executive guide

What AI can really do for a company in 2026, where the value is, why most efforts stall, and how to start in the next 90 days.

ABOUT THIS GUIDE

A plain-language briefing for the people who decide where the money goes.

Almost every company now uses AI somewhere. Very few can point to it in their profit and loss. This guide explains why, using the most recent global research, and turns it into decisions a Lebanese executive team can take this quarter. It assumes no technical background.

01	Where enterprise AI stands	Wide use, narrow value
02	What AI actually does today	Four kinds of AI, in plain terms
03	Agentic AI without the hype	What agents are and where they fail
04	Where the value is	Functions, sectors and the back office
05	Why most AI efforts stall	Five patterns to avoid
06	What the leaders do differently	Six habits from the high performers
07	The Lebanese context	Constraints, advantages and the opening
08	Your first 90 days	A plan, and questions for any vendor

READ THIS IF

You sit on an executive committee or board and keep being asked to approve AI spending without a clear way to judge it.

TIME TO READ

About 35 minutes. Each chapter opens with its main point, so you can skim the bold lines first.

PART OF A SERIES

No. 02 covers readiness, No. 03 the path from pilot to production, No. 04 governance and data.

CHAPTER 01

Where enterprise AI stands

AI is now used almost everywhere, but it pays off in a small minority of companies. That gap, not the technology, is the story of 2026, and it is the opening for any company willing to be deliberate.

Three years after generative AI arrived in offices, adoption is no longer the question. McKinsey's 2026 global survey found that nearly nine in ten organisations regularly use AI in at least one business function, and that more of them now use it in three or more functions than a year earlier.[1] Stanford's AI Index puts organisational adoption at 88% for 2025, up from 78% the year before.[8]

Value is a different matter. In the same McKinsey survey, only 37% of respondents attribute any impact on earnings (EBIT) to AI, a share that did not move from the previous year. The group McKinsey calls "AI high performers", companies that get 5% or more of their EBIT from AI, is about 6% of respondents.[1]

From using AI to earning from it

Share of organisations, McKinsey Global Survey on AI, 2026



SOURCE: MCKINSEY, THE STATE OF AI IN 2026 (1,719 RESPONDENTS, 97 COUNTRIES)

Other research draws the same line. Boston Consulting Group's 2026 index of 1,330 senior leaders classes 7.5% of companies as "future-built", the most AI-mature, and a further 41% as scaling. Compared with laggards, the future-built group delivers 2.3 times the total shareholder return, 2.4 times the revenue growth and 2.8 times the EBITDA growth.[3] A widely cited 2025 study from MIT's Project NANDA, based on 150 leader interviews, 350 employee surveys and 300 public deployments, found that about 95% of enterprise generative AI pilots produced no measurable profit and loss impact.[5]

WHAT THIS MEANS FOR YOU

The technology is available to everyone at roughly the same price. What separates the winners is how they choose problems, prepare data, change the work and measure the result. Those are management capabilities, and they can be built.

The good news hidden in the numbers

BCG's headline for 2026 is that AI "is starting to pay off": future-built and scaling companies together now make up nearly half the market, and the scaling group alone earns 1.8 times the shareholder return of laggards.[3] Value is no longer confined to a handful of global giants. Companies that move from experiments

to a few well-chosen, well-run use cases are catching up quickly.

Spending is also moving. BCG reports that AI budgets roughly doubled, from about 1.7% of revenue in late 2025 to 3.3%, and that more than 80% of AI spending now sits outside the IT budget, in the business units themselves.[3] AI has become an operating decision, not an IT project, which is exactly why executives need their own view of it.

Four things executives still hear, and what the evidence says

WHAT YOU HEAR	WHAT THE EVIDENCE SAYS
"We need to wait until AI is mature."	Capability rose sharply in 2025 and prices keep falling. The companies pulling ahead started with narrow, practical cases and reinvested what they learned.[3,8]
"Buying licences for everyone is our AI strategy."	Individual productivity rises, but EBIT impact comes from redesigning specific workflows. Tools alone rarely show up in the accounts.[1]
"Most AI projects fail, so it is a gamble."	Failure is concentrated in projects with no business owner, no data plan and no road to production. Those are avoidable.[5,7]
"AI will replace our workforce."	In 2026, 14% of respondents said AI had reduced their total headcount over the past year, while a year earlier 32% had expected declines. Roles are changing faster than they are disappearing.[1]

CHAPTER 02

What AI actually does today

"AI" covers four quite different tools. Knowing which one a proposal relies on tells you most of what you need about its cost, its risk and how quickly it can work.

KIND OF AI	WHAT IT DOES	TYPICAL ENTERPRISE USE	WATCH FOR
Predictive AI	Learns patterns from your historical data to forecast, score or classify.	Demand forecasting, credit scoring, churn risk, fraud flags, predictive maintenance.	Needs clean, structured history; drifts when conditions change.
Generative AI	Reads and writes text, images and code from instructions in plain language.	Drafting reports and replies, summarising documents, translating, coding assistance.	Confident errors ("hallucinations"); output needs review.
AI on your documents	Generative AI connected to your own files and systems so answers cite your sources. Often called retrieval, or RAG.	Policy and procedure assistants, contract search, technical manuals, customer service knowledge.	Only as good as the documents behind it; access rights must carry over.
Agents	AI that plans steps and uses tools (email, ERP, browser) to complete a task with limited supervision.	Processing invoices end to end, preparing a credit file, triaging and routing requests.	Hardest to control; needs permissions, logs and approval points.

What it does well, and where it fails

Modern generative models are strong at language-heavy work: reading long documents, extracting information, drafting in a given tone, switching between Arabic, French and English, and writing code. Stanford's 2026 AI Index notes that performance on a key software-engineering benchmark rose from about 60% to near 100% in a single year.[8]

They are weak where a single wrong answer is expensive and nobody checks: precise arithmetic without a tool, facts outside the material they were given, and decisions that depend on context nobody wrote down. They can also be inconsistent, answering the same question differently on different days. Good enterprise systems are designed around these weaknesses with source citations, structured outputs, validation rules and human review at the points that matter.

280×

Fall in the price of GPT-3.5-level AI between November 2022 and October 2024, from about $20 to $0.07 per million tokens.

STANFORD AI INDEX 2025

~47%

Average fall per quarter in the cost of a given level of AI performance over the past three years.

EPOCH AI, SEPT 2026

25

Elo points separating six leading model makers on a public leaderboard in March 2026.

STANFORD AI INDEX 2026

Two trends matter for budgeting. First, the cost of a given level of capability keeps collapsing.[9,10] Second, the leading models have converged in quality, so competition has shifted to cost and reliability.[8] For an executive, the conclusion is practical: **the choice of model is rarely the decision that makes or breaks a project.** The workflow, the data and the controls around the model are.

> **A caution on cheap tokens.** Lower unit prices do not guarantee lower bills. In McKinsey's 2026 survey, about one in five organisations said AI operating costs, including token costs, already constrain their use of AI.[1] Usage grows as prices fall, so cost tracking belongs in every business case.

A short glossary for the boardroom

TERM	WHAT IT MEANS IN PRACTICE
Model	The trained AI engine, such as a large language model. Usually rented from a provider, sometimes run on your own servers.
Token	The unit AI providers charge by, roughly three quarters of an English word. Arabic text often uses more tokens per word.
Prompt and context	The instructions and material given to the model for a task. Better context, such as the right documents, matters more than clever wording.
Hallucination	A fluent but false answer. Reduced by giving the model your sources, asking it to cite them, and checking outputs.
Retrieval (RAG)	Searching your own documents for the relevant passages and handing them to the model, so answers are grounded in your material.
Fine-tuning	Further training a model on your examples to fix a style or format. Rarely the first step; good context usually does more, for less.
Open-weight model	A model whose files can be downloaded and run on your own infrastructure, useful when data cannot leave the premises.
Evaluation	A fixed set of real test cases with known right answers, used to check quality before launch and after every change.

CHAPTER 03

Agentic AI without the hype

Agents are the most promising and the most oversold part of AI right now. Treat autonomy as something a process earns step by step, not something you buy.

An agent is an AI system that does more than answer: it decides which steps to take and carries them out through other software, such as reading an inbox, updating a record or sending a message. That is genuinely new, and the money is following it. BCG estimates that agents already produced about 17% of total AI value in 2025 and expects agentic AI to account for around 40% of all AI value by 2030.[4,3]

The failure rate is also new. Gartner predicts that more than 40% of agentic AI projects will be cancelled by the end of 2027 because of escalating costs, unclear business value or inadequate risk controls. It also warns of "agent washing", the relabelling of ordinary chatbots as agents, and estimates that only about 130 of the thousands of vendors claiming agentic products are genuine.[6]

40%+

of agentic AI projects expected to be cancelled by the end of 2027.

GARTNER, JUNE 2025

42%

of companies expect to give AI agents real decision-making authority by 2030.

BCG, SEPT 2026

5%

have the full set of controls in place for that today.

BCG, SEPT 2026

The autonomy ladder

A useful way to judge any agent proposal is to ask which rung it sits on, and whether the controls match.

RUNG 1

Assist

AI drafts or suggests; a person does the work and decides. Example: a draft reply to a supplier.

RUNG 2

Automate a step

AI completes one defined step inside a fixed process. Example: extracting fields from invoices.

RUNG 3

Run a workflow

AI moves a case through several steps, with a person approving at set points. Example: preparing a claim file.

RUNG 4

Act on its own

AI decides and acts within limits, and people review samples and exceptions. Example: routing and resolving routine tickets.

Most companies should spend their first year on rungs 1 to 3. Each rung up requires more of the same three things: **clear permissions** (what the agent may touch), **a full record** of what it did and why, and **defined points where a person approves** or can stop it. Gartner's own advice is to pursue agents only where they deliver clear value or return, and notes that many cases marketed as agentic do not need an agent at all.[6]

FIVE QUESTIONS BEFORE APPROVING AN AGENT

- Could a simpler fixed workflow, with AI in one or two steps, do the same job more reliably?
- What is the worst thing it could do with the permissions it is asking for, and who would notice?
- Where does a person approve, and what happens when the agent is unsure?
- Is every action logged in a way an auditor could follow?
- What does it cost per completed case, including failed attempts and human review?

Worked example: supplier invoices at a distributor

A distributor receives about 3,000 supplier invoices a month by email, as PDFs and phone photos, in three languages. Two clerks key them into the ERP and match them to purchase orders.

RUNG	WHAT THE AI DOES	CONTROLS THAT MAKE IT SAFE
2	Reads each invoice and fills a draft entry: supplier, date, lines, totals, tax.	Clerk confirms every entry; low-confidence fields highlighted.
3	Also matches the invoice to the purchase order and goods received, and flags mismatches.	Clerk approves matches; any mismatch above a threshold goes to a supervisor.
4	Posts fully matched invoices under a set amount without review.	Daily sample audit, amount limits, full log, one-click reversal, monthly accuracy report.

Most companies capture the bulk of the value at rung 3, where clerks move from typing to checking exceptions. Rung 4 is worth it only once accuracy has been measured on several months of real volume.

CHAPTER 04

Where the value is

Companies tend to spend on the visible front office. The fastest, most measurable returns are often in the back office, where work is repetitive, document-heavy and already costed.

McKinsey's respondents most often report revenue gains from AI in marketing and sales and in product or service development.[1] MIT's research adds an important twist: more than half of generative AI budgets go to sales and marketing tools, yet the clearest returns appeared in back-office automation, such as replacing outsourced processing, cutting external agency spend and streamlining operations. Companies that crossed the divide reported savings of $2 million to $10 million a year from reduced outsourcing and a 30% cut in agency spend.[5]

RULE OF THUMB

Look first where people read, re-type, check or route documents all day. The work is already measured, the savings are easy to prove, and the risk of a customer seeing a mistake is lower.

Where Lebanese companies typically start

The examples below are illustrative starting points drawn from common patterns in each sector. Your own list should come from your operations, scored with the method in No. 03 of this series.

BANKING AND INSURANCE

Documents and compliance

KYC file review, credit memo drafting, claims intake and triage, regulatory reporting support, multilingual customer correspondence.

MANUFACTURING

Quality and uptime

Maintenance logs turned into failure predictions, visual quality inspection, production planning assistants, supplier document checks.

DISTRIBUTION, FMCG, RETAIL

Orders and demand

Order capture from WhatsApp and email into the ERP, demand forecasting by outlet, route planning, promotion analysis.

HEALTHCARE

Administration first

Insurance claims coding and checks, discharge summary drafting, appointment and call-centre handling, stock forecasting.

LOGISTICS AND PORTS

Paperwork flow

Customs and shipping document extraction, exception handling, shipment status queries, capacity forecasting.

PROFESSIONAL SERVICES

Knowledge work

Proposal and report drafting, contract review, research summaries, time and billing analysis.

Three kinds of value, and how to measure each

VALUE	WHAT CHANGES	MEASURE IT BY
Efficiency	The same work takes fewer hours, less outsourcing or fewer errors.	Cost per case, hours saved and actually redeployed, error and rework rates.
Growth	Faster responses, better targeting or new services bring in revenue.	Conversion, response time, revenue per customer, new revenue lines.
Resilience	Knowledge stays in the company; decisions are faster and more consistent.	Time to decision, dependency on key individuals, audit findings.

McKinsey found that 80% of respondents say AI has improved their individual productivity, yet only 37% see an effect on EBIT.[1] The difference is usually that saved time is never converted into lower cost or higher output. **Decide before a project starts what will happen to the time it frees.**

A QUICK WAY TO SIZE AN EFFICIENCY CASE

Annual value = cases per year × minutes saved per case ÷ 60 × loaded cost per hour × share of saved time actually redeployed − annual running cost

Example: 36,000 invoices a year, 6 minutes saved on each, $9 loaded cost per hour, 70% of the time redeployed, $6,000 a year to run. That is 36,000 × 6 ÷ 60 × 9 × 0.7 − 6,000 = **about $16,700 a year** from one process, before counting fewer errors and faster supplier payments. If the answer is small, the case is either too narrow or should be judged on growth or risk instead.

CHAPTER 05

Why most AI efforts stall

Projects rarely fail because the AI is not clever enough. They fail for five predictable, management-level reasons.

1 **A technology looking for a problem.**
The project starts from a tool or a vendor demo rather than a costed business problem. Gartner lists poor use-case selection and unclear business value as the most common reasons generative AI projects are abandoned; it found at least half were dropped after proof of concept by the end of 2025.[7]

2 **A pilot with no road to production.**
The pilot is built as a demonstration, so nothing about integration, security, support or cost has been solved when it "succeeds". MIT found only about 5% of custom enterprise AI tools reach production, and that large enterprises take nine months or more to move from pilot to full implementation, against roughly 90 days for mid-market firms.[5]

3 **Data that is not ready.**
The information exists, but in paper files, personal inboxes and spreadsheets nobody trusts. Gartner predicts that through 2026 organisations will abandon 60% of AI projects that are not supported by AI-ready data, and that 63% of organisations lack, or are unsure they have, the right data practices.[7]

4 **Nobody owns the change.**
The system works, but the team keeps doing things the old way. BCG's rule of thumb is that successful AI programmes put 10% of their effort into algorithms, 20% into technology and data, and 70% into people and processes.[12]

5 **Ungoverned use.**
Staff use personal AI accounts with company data because the official tools are slow to arrive. MIT reports that over 90% of employees in its sample used personal AI tools for work, while only about 40% of companies had bought official subscriptions.[5] IBM found that breaches at organisations with high levels of such "shadow AI" cost an average of $670,000 more.[11]

Build, buy or partner?

One of MIT's clearest findings concerns sourcing. AI tools bought from specialised vendors or built with external partners reached deployment about 67% of the time, against roughly 33% for purely internal builds.[5] The lesson is not "never build". It is that teams new to AI underestimate integration, evaluation and change work, and benefit from people who have done it before. A sensible pattern is to partner for the first projects while deliberately building internal capability to own and run what is delivered.

> **Our view at HTech.** Every one of these five patterns shows up before a line of code is written. That is why NEAP starts with readiness and case selection, and why the first deliverable of a good AI engagement is a decision, not a demo.

CHAPTER 06

What the leaders do differently

The companies that earn from AI are not using different technology. They manage it differently, and the research is consistent about how.

High performers versus everyone else

Share of respondents who fundamentally redesign workflows around AI, McKinsey 2026

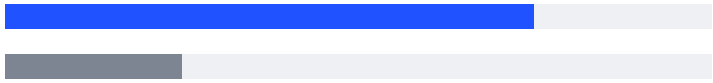

SOURCE: MCKINSEY, THE STATE OF AI IN 2026

McKinsey's high performers are about three times as likely as others to redesign workflows around AI rather than adding AI to existing steps. They are twice as likely to report committed senior leaders and defined processes for measuring AI's impact, and more than twice as likely to spend over 15% of their technology budget on AI.[1] BCG describes the same pattern as a formula: strategic clarity on where AI creates value, applied AI that redesigns work end to end, investment in governance and people alongside technology, and impact measured directly in the profit and loss.[3]

Six habits you can copy

01

Pick few, go deep

Concentrate on a short list of use cases tied to the company's biggest priorities, instead of dozens of scattered pilots.

02

Redesign the work

Ask how the process should run if AI did the reading and drafting, then rebuild it. Do not just speed up the old steps.

03

Lead from the top

A named executive owns each use case and its result. Senior leaders use the tools themselves and are seen to.

04

Measure in money

Set a baseline before starting, track the metric that matters, and report the effect on cost or revenue, not usage.

05

Govern early

Clear rules on data, tools and approvals from day one make it safe to move faster, not slower.

06

Build people

Train widely and create AI roles. BCG reports that seven in ten future-built and scaling companies are already retraining staff.[3]

> **The leadership test.** Ask your team three questions: which three AI cases matter most to us, who owns each one, and how much money has each made or saved so far? If the answers take longer than a minute, the habits are not in place yet.

CHAPTER 07

The Lebanese context

Lebanese companies face real constraints, but several of the factors that decide success with AI play in their favour.

The regional prize is large. PwC estimates that AI could contribute about $320 billion to the Middle East economy by 2030, roughly 11% of the region's GDP.[13] At home, Lebanon now has a Ministry of State for Technology and AI, whose five-year LEAP strategy aims, among other goals, to place Lebanon among the world's top 50 AI users by 2030.[14] Public strategy, however, mainly addresses the state and the startup ecosystem. The established private sector, the banks, manufacturers, distributors and hospitals that employ most people, has to move on its own initiative. That is the gap NEAP was created to close.

CONSTRAINTS TO PLAN FOR

What makes it harder

- Information held on paper, in WhatsApp threads and in personal spreadsheets.
- Documents in Arabic, French and English, often mixed.
- Unreliable power and connectivity at some sites.
- Tight capital budgets and short payback expectations.
- Data protection rules (Law 81/2018) that many firms have not yet mapped.

ADVANTAGES TO USE

What makes it easier

- Mid-sized companies move from pilot to production faster than large ones.[5]
- Falling AI prices make small, focused projects affordable.[9]
- Strong universities and a trilingual talent pool.
- Short decision chains in family and founder-led firms.
- Plenty of manual, document-heavy work where returns are easiest to prove.

Designing for local conditions

A few design choices go a long way. Choose tools that handle Arabic well and test them on your real documents, including scanned and handwritten ones. Prefer solutions that keep working when connections drop, or that run on your own servers when data cannot leave the premises. Start with processes whose savings show up within a quarter. And put a simple acceptable-use policy in place before staff paste client data into public tools.

HOW NEAP HELPS

NEAP brings companies through five phases: a shared foundation in AI for leadership, an honest readiness assessment, a short list of cases chosen by feasibility and consequence, implementation inside operations with NEAP Fellows from partner universities on the team, and finally the training that lets the company run on its own.

CHAPTER 08

Your first 90 days

You do not need an AI strategy document to start. You need an owner, a few rules, an honest baseline and one well-chosen case run with production in mind.

WEEKS 1–2

Set up

Name an executive owner. Publish a one-page acceptable-use policy. Take the NEAP Readiness Check with your leadership team and compare answers.

WEEKS 3–6

Choose

List 10–20 candidate cases from real operations. Score them on value and feasibility. Check the data behind the top three. Pick one, at most two.

WEEKS 7–12

Prove

Run a pilot built to go live: real users, real data, a baseline, a target, a cost per case and a go or stop decision on a fixed date.

Questions to ask any AI vendor or consultant

- [ ] Which of our business metrics will this change, by how much, and how will we measure it against today's baseline?
- [ ] What does it cost per transaction or per user once it runs at full volume, including model usage, hosting and support?
- [ ] Where does our data go, who can see it, is it used to train anyone's models, and can it stay in Lebanon or on our servers?
- [ ] How does it handle Arabic and mixed-language documents? Can we test it on our own samples before signing?
- [ ] What happens when the AI is wrong or unsure? Where does a person review, and how are errors caught?
- [ ] How does it connect to our existing systems, and who maintains those connections?
- [ ] Which reference clients run this in production today, not in a pilot?
- [ ] How will our own team learn to run, adjust and own it after you leave?

Continue with the series. No. 02 helps you score your readiness across five pillars. No. 03 gives the full method for choosing and scaling cases. No. 04 sets out the governance and data rules to put in place first.

SOURCES AND NOTES

Where the numbers come from

1. McKinsey & Company, *The state of AI in 2026: On the road to ROI*, McKinsey Global Survey, published 25 August 2026 (fielded 4 May–8 June 2026; 1,719 participants in 97 countries). mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
2. McKinsey & Company, *The state of AI in 2025: Agents, innovation, and transformation*, 2025 (1,993 participants, fielded June–July 2025).
3. Boston Consulting Group, *AI Is Starting to Pay Off: Almost 50% of Companies Now Generate Value with It* (Applied AI Index 2026), press release, 30 September 2026; survey of 1,330 CxOs and senior leaders.
4. Boston Consulting Group, *The Widening AI Value Gap: Build for the Future 2025*, September 2025.
5. Challapally, A., Pease, C., Raskar, R. and Chari, P., *The GenAI Divide: State of AI in Business 2025*, MIT Project NANDA, July 2025; as reported by Fortune (18 August 2025) and others.
6. Gartner, *Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027*, press release, 25 June 2025.
7. Gartner, *Lack of AI-Ready Data Puts AI Projects at Risk*, 26 February 2025; and Gartner, *Why 50% of GenAI Projects Fail*, gartner.com/en/articles/genai-project-failure.
8. Stanford Institute for Human-Centered AI, *AI Index Report 2026*, 13 April 2026.
9. Stanford Institute for Human-Centered AI, *AI Index Report 2025*: inference cost for GPT-3.5-level performance, November 2022 to October 2024.
10. Epoch AI, *The plunging price of thought*, 22 September 2026.
11. IBM and Ponemon Institute, *Cost of a Data Breach Report 2025: The AI Oversight Gap*, 30 July 2025 (600 organisations).
12. Boston Consulting Group, AI at Scale and the 10-20-70 rule, bcg.com/capabilities/digital-technology-data/data-analytics.
13. PwC Middle East, *US$320 billion by 2030? The potential impact of Artificial Intelligence in the Middle East*, 2018.
14. L'Orient Today, reporting on the Ministry of State for Technology and AI's LEAP strategy (Launch, Apply, Advance, Promote), 2025.

This guide is published by HTech S.A.R.L. as part of the National Enterprise AI Program (NEAP), an HTech program. It is general information for business leaders, not legal, financial or regulatory advice; check how any law or standard applies to your company with qualified counsel. Figures are quoted from the sources listed, as published by them, and were checked in November 2026. You are welcome to share this guide unchanged and free of charge. © 2026 HTech S.A.R.L.

YOUR NEXT STEP

Find out where your company stands before you spend on AI.

Every company in this guide that got value from AI started from an honest picture of its own readiness. The NEAP Readiness Check gives you that picture in three minutes, and a free debrief with HTech turns it into a short list of first moves.

FREE · 3 MINUTES

Take the Readiness Check

Ten questions across the five pillars of the NEAP Readiness Index. You get your level, your weakest pillar and the one next step that fits.

FREE · 30 MINUTES

Book a debrief with HTech

Walk through your result with an HTech consultant and leave with a short list of where AI can create value in your company first.

Book a debrief

NEXT NEAP COHORT

Join the Foundation Workshop

One day for your leadership team. You leave with a full readiness profile and a commitment card for your first AI case.

joinneap.com

HTech S.A.R.L. · Berytech, Mathaf, Beirut, Lebanon
info@htechco.com · +961 81 733 831 · htechco.com · joinneap.com

HTech / National Enterprise AI Program™