



34 CASES

7 REMAIN

— INSIGHT · STRATEGY & USE CASES

The Use Cases HandsOn Turns Down.

A strategy sets the guardrails — it selects what will not be done before it selects what will. The refusal set is the first movement of any serious AI strategy.

AUTHOR

SERIES

EDITION

PUBLISHED

READ TIME

HandsOn Insights Pillar 01 · Strategy & Use Cases Vol. 01 · N° 07 12 May 2026 5 minutes

EXECUTIVE SUMMARY

An AI portfolio is not yet a strategy.

Five findings the board should weigh before the next use-case list.

01**Strategy means deletion first, selection second.**

A strategy selects, on the basis of defined guardrails, what will *not* be done before it selects what will. The portfolio is nothing more than an initial shortlist — by no means a strategy.

02**25–30 %. The rest is quietly absorbed into the line.**

Only a narrow third of mid-market AI initiatives move from pilot status into productive operation. The rest are maintained over subsequent quarters, occasionally renamed — and written off at year-end.

03**Three filters separate a use case from a distraction.**

Operating-Model Fit. Cost of Autonomy. Strategic Congruence. What does not survive these three filters does not belong in the portfolio — it belongs in a different discussion, or in the bin.

04**The real portfolio is small. In our example: 7 of 34.**

Four cases carry the strategic core; three are learning cases — chosen to build internal capability, not for ROI. The other 27 are not deferred. They are deleted.

05**Three anti-patterns that look like strategy — and aren't.**

The matrix with circles. The sequencing lie. The quick-win reflex. All three dodge the hardest part — the strategic choice. No one says no. So no one has decided.

CONTENTS

What you'll find here.

01	A portfolio is not a strategy	04
	The fictional industrial supplier, €1.2 bn revenue, 34 use cases. How a list of sponsors, business cases and traffic lights gets relabelled as strategy.	
02	What the market does with use-case strategy	07
	The 2×2 matrix with circles. The "we'll do all of them" reflex. No one says no — so no one has decided.	
03	How HandsOn does use-case selection	10
	Refusal-first. Two movements, three filters. What remains is 7 of 34 cases — and a real portfolio.	
04	What good looks like — and three anti-patterns	13
	Three signatures of a genuinely strategic use-case portfolio. The three reflexes that destroy it.	
05	When a Strategy & Use Cases engagement fits	16
	Five signals — three or more hits, and use-case selection is the right door.	
06	About HandsOn & engagement	19
	Method, sources, contact — and the AI Strategy Engagement.	

SECTION

01

A portfolio is not a strategy.

"We have an AI strategy portfolio now." We often hear this from executives. A strategy is something else: it selects what will not be done before it selects what will.

34 cases, 31 on green, nine months of defending.



On the management team's desk of a fictional client — industrial supplier, €1.2 bn revenue, three plants, one division losing margin for two years — sits a list of 34 AI use cases. Each case has a sponsor, an estimated business case, and a traffic light. Thirty-one of them are green.

The head of strategy calls the document the "AI portfolio." The HandsOn response to such a list is *not* to help with prioritisation. We start by deleting use cases.

After three hours of methodical work on the refusal set, seven cases remain. Twenty-seven drop out. This usually leads directly to discussions with the head of strategy, who painstakingly compiled the list in the first place. He is taken aback — because he has defended the document for nine months.

And the pattern is set, the one we see in every other board meeting: an AI project portfolio and an AI strategy are not the same thing. The portfolio is an initial shortlist of potential use cases — but by no means a strategy.

A strategy differs in that it selects, on the basis of defined guardrails, **what will not be done** before it selects what will. That very movement is missing from almost every AI portfolio we get to look at.

"No one says no. So no one has decided."

HANDSON · OBSERVATION FROM BOARDROOMS

25–30 %. The rest is quietly absorbed into the line.

BCG/MIT Sloan · *State of AI 2024* puts the share of mid-market AI initiatives that move from pilot status into productive operation in the range of 25 to 30 percent. The rest are maintained over subsequent quarters, occasionally renamed, and at year-end quietly absorbed into the line.

25–30%

Reach productive operation.

BCG/MIT SLOAN · STATE OF AI 2024

70%

Stay in pilot or are absorbed into the line.

BCG/MIT SLOAN · STATE OF AI 2024

18 mo.

Typical roadmap stretch when every case stays in.

HANDSON · MID-MARKET OBSERVATION

WHAT'S LEFT AT YEAR-END

The uncomfortable question: what did implementing those 34 cases actually accomplish?

Cases that didn't get implemented. Or — far worse — cases that got implemented without measurable ROI and without adoption by the organisation. Activity and a long list of cases to be implemented provide a sense of security. At least until the year-end tally.

SECTION

02

The matrix with circles. The second reflex: "We'll do all of them."

Two exercises that look like strategy — and aren't. They produce a document that makes the board feel good, and a pilot portfolio that subsequently fails to scale.

The second reflex: "We'll do all of them, because everything is important."

The portfolio remains at 34 cases, the roadmap stretches to 18 months. The scarcity inherent in any strategy gets quietly glossed over: scarcity of attention, of leadership bandwidth, of data architecture, of employee adoption capacity.

-
- **Attention.** The board cannot hold 34 cases in mind at the same time. At twelve it gets tight. At seven, every case can be named without a cheat sheet.

 - **Leadership bandwidth.** Every case needs a sponsor to carry it through the line. In Q3 the board realises that the bandwidth is enough for 4 cases at once, not 12.

 - **Data architecture.** Every second case quietly waits on a data source that doesn't yet exist, or isn't yet accessible. That only surfaces in month six.

 - **Adoption capacity.** Employees cannot integrate eight new tools into their work in parallel. Try it, and adoption falls to near zero — at the full implementation bill.

SO NO ONE HAS DECIDED

The portfolio looks like a strategy, but isn't.

No one says no. No one has lost anything. No one has stuck their neck out. Eight of the twelve fall quiet — in Q3, when bandwidth gets tight. Whoever had decided that in Q1 would have saved nine months of attention.

SECTION

03

Refusal-first. Deletion first. Selection second.

The HandsOn method has two movements. The first is deletion. Selection comes second. Before a single case lands on a matrix, we write down which classes of applications the company will *not* pursue this fiscal year.

Movement one: the refusal set. Three filters.

Before a single case lands on a matrix, we write down which classes of applications the company will not pursue this fiscal year. We use three filters — each one is hard, and on its own sufficient to lift a case out of the portfolio.

FILTER 01 · REFUSAL

01

Operating-Model Fit

Which cases would require a steering logic the company does not have today? A case that would re-cut decision rights across three functional areas, without an instance to decide that cut, does not belong in the portfolio. **It belongs in the discussion about the operating model** — not in the next implementation wave.

FILTER 02 · REFUSAL

02

Cost of Autonomy

Every case has an autonomy level: suggestion, decision with human in the loop, decision without, autonomous action. Each step up raises governance costs disproportionately. **When control eats the business case, the case scales back to Level 1 — or it is dropped.**

FILTER 03 · REFUSAL

03

Strategic Congruence

What does the business model turn on? A supplier whose margin depends on time-to-quote should not spend energy on an HR chatbot. **Cases that do not pay into the one or two strategic levers of the business model are distractions** — even when they work.

METHOD

The three filters are **OR-joined**, not **AND**.

A case must pass all three filters to stay in the portfolio. Fail just one, and the clean answer is deletion — not deferral to a later wave. Deferral is the sequencing lie in another form.

Movement two: selection within the remainder.

What survives the three filters is the real portfolio. Applied to the fictional industrial supplier: 7 of 34 cases. Four of them carry the strategic core; three are learning cases — chosen to build internal capability, not for ROI. The other 27 are not "deferred." They are deleted.

4

Strategic core — cases that pay into the one or two strategic levers.

EXAMPLE CASE · 34 → 7

3

Learning cases — chosen to build internal capability, not for ROI.

EXAMPLE CASE · 34 → 7

27

Deleted. Not deferred.

EXAMPLE CASE · 34 → 7

Only here does matrix logic come into play, and it is tighter than usual. We score the remaining cases against the **HandsOn AI Operating Model** — six domains, from Strategy to Operations — and check, for each case, in which domain the bottleneck sits.

BOTTLENECK LOGIC

When the bottleneck is not in Strategy or Use Case, but in Governance or Literacy — the answer is: not yet.

Address the bottlenecked domain first, then start the case. Try it the other way around, and you're building on a foundation that catches up with you in month six — usually in the form of near-zero adoption or regulatory exposure that nobody had on the radar.

SECTION

04

Three signatures. Three anti-patterns. One slide.

A use-case portfolio after the refusal-first cut is recognisable by three things. It gets destroyed by three reflexes we see in every other meeting.

Three signatures of a genuinely strategic use-case portfolio.

SIGNATURE 01

7

Small. Seven, not 34.

The board can name each case without a cheat sheet. Seven cases, not thirty-four. Anyone who can't explain the seventh from a standing start hasn't understood the seventh — and should delete it.

SIGNATURE 02

One slide

The refusal set is in the document, visibly.

On its own slide stands which application classes the company will not address this fiscal year, and why. It's the slide the CFO asks for next quarter — and the supervisory board the quarter after.

SIGNATURE 03

L1

Most cases start at Level 1.

Level 1 = suggestion with human in the loop. The path upward is explicit, condition-bound, anchored in the governance chapter. No one drifts quietly from L1 to L3 — that is the most common source of regulatory exposure.

These three signatures are not a wish list — they are the litmus test. Miss one of them, and the portfolio isn't the result of a strategic decision, but of a distribution.

"Seven cases. A refusal slide. Level 1 as the default. If you can't show that, you don't have a strategy."

HANDSON · REFUSAL-FIRST TEST

Three reflexes that look like strategy — and aren't.

Anti-Pattern 01

MARKET
STANDARD

The matrix with circles. 2x2 axes, circle size as investment. Produces a document that makes the board feel good, and a pilot portfolio that subsequently fails to scale. The hard cases — usually the ones with the highest structural impact — drift into "next quarter."

Anti-Pattern 02

MOST COMMON
FORM

The sequencing lie. All 34 cases stay in, they are just distributed across 18 months. No one has decided, no one has lost anything, no one has stuck their neck out. In Q3 the board realises the bandwidth is enough for 4 cases at once, not 12. The other 8 fall quiet.

Anti-Pattern 03

MOST
DANGEROUS
FORM

The quick-win reflex. An AI chatbot for HR self-service produces a nice press release and zero cents of margin impact. A use-case portfolio that consists of three quick wins and a roadmap has dodged the hardest part — the strategic choice.

COMMON DENOMINATOR

All three anti-patterns avoid the uncomfortable step: disappointing someone.

Deletion means: a sponsor loses their case. A function loses its budget. A slide gets pulled from the board deck. Every anti-pattern is an elegant way to avoid that disappointment — at the cost of nine months of attention at year-end.

SECTION

05

Five signals. Three hits. The right door.

Not every company that needs AI advisory needs use-case strategy first. Five signals make the difference — three or more hits, and use-case selection is the right door.

Five signals for the right engagement moment.

SIGNAL**01****The list on the table.**

The board has a list of more than 15 cases in front of it and no shared logic for which five of them count this year.

SIGNAL**02****Parallel initiatives.**

In the last twelve months there have been multiple divisional initiatives with their own budgets — Marketing, HR, Operations — running in parallel and not talking to each other.

SIGNAL**03****CFO veto.**

The CFO returned the last investment request with the question of what happens if all three cases come good at the same time.

SIGNAL**04****No shared term.**

There is no shared term for what a "strategic AI case" is at this company, and what is mere tool delivery.

SIGNAL**05****External obligation.**

An external obligation — EU AI Act, results call, supervisory board — demands a position, and the current answer is "we're working on it."

When the bottleneck sits elsewhere — no inventory of running applications, no clarity on high-risk classification, a capability gap that hollows out any strategy — the strategy engagement is not the first movement. The comprehensive **HandsOn Diagnostic** leads, to identify the exact pain points first.

27 cuts. One slide. The longest discussion of the year.

Three weeks after the cuts — again, in the fictional example — the head of strategy presents the refusal set to the board. There are 27 cuts listed on a single slide. The discussion in that meeting is the longest of the year.

Most board members find it difficult to prioritise. Activity and a long list of cases to be implemented provide a sense of security. At least until the year-end tally.

That's when it becomes clear that cases were not implemented — or, much worse, were implemented without measurable ROI and without adoption by the organisation. Then you have to face the far more uncomfortable question: **what did implementing those 34 cases actually accomplish?**

That's why at HandsOn we always start by trimming the use-case list, even when it means things get uncomfortable right away. The discomfort belongs at the start — not at the end of the year, when it can no longer be corrected.

"Activity and a long list give a sense of security — until the year-end tally."

CLOSE · THE USE CASES HANDSON TURNS DOWN

STRATEGY & USE CASES · ENGAGEMENT

Which cases would you delete?

The HandsOn AI Strategy Engagement helps you identify the right use cases and then move on to focused implementation — beginning with the refusal set, three filters, and the honest question about the strategic levers of your business model.

ABOUT THE PUBLISHER

HANDSON

HandsOn Insights.

Pillar 01 · Strategy & Use
Cases

HandsOn is an independent advisory firm working with European mid-market and enterprise leadership teams on the organisational design of AI. The Insight series publishes proprietary research and field-tested frameworks on the strategy, structure and governance of AI in large organisations.

This whitepaper describes the HandsOn Refusal-First method for use-case selection — three filters, two movements, one real portfolio. Applied to a fictional industrial supplier: 34 cases on the table, 7 after three hours of methodical work.

STRATEGY & USE CASES · ENGAGEMENT

**Which cases would you delete?
Request the AI Strategy
Engagement.**

[BOOK A MEETING →](#)

SOURCES

- **BCG/MIT Sloan** · State of AI 2024 · 25–30 % pilot-to-production rate in the mid-market.
- **HandsOn** · Refusal-First method · three filters, two movements.
- **HandsOn AI Operating Model** · six domains · Foundation and Activation Layer.
- **HandsOn** · Cost-of-Autonomy logic · four autonomy levels · Whitepaper Vol. 01 · N° 05.
- **EU AI Act** · Regulation (EU) 2024/1689 · high-risk classification.
- **Example case** · fictional industrial supplier · €1.2 bn revenue · 34 → 7 cases.

CONTACT

info@wearehandson.de
wearehandson.de

SERIES

HandsOn Insight
Pillar 01 · Strategy & Use
Cases

EDITION

Vol. 01 / N° 07
12 May 2026

© 2026 HandsOn. All rights reserved. This publication contains general information. With this publication, HandsOn does not provide business, financial, legal, tax or other professional advice. Statistics remain the property of their respective publishers. The example case (industrial supplier, €1.2 bn revenue) is fictional.