

The request is not the problem

The Performance Consultant AI: how a training request leads to a grounded decision about performance, in minutes rather than weeks.

Built on human performance improvement

Gilbert, Mager and Pipe, Rummler and Brache, and the Dutch line of Arets, Overduin and Heijnen. Half a century of method, put within reach of every learning advisor.

Diagnosis

retail · 1,000 people · request: customer-experience training

DIAGNOSTIC CONFIDENCE
moderate to high

THE GAP

sales per visit €18.40, target €21.00, worth about €2.6M a year

WHAT IS CAUSING IT

- out-of-stocks on the top 40 lines process · triangulated
- pricing errors at the register tools · substantiated
- no feedback on visit value per store expectations · hypothesis

RECOMMENDED NEXT STEPS

fix replenishment and the price file first; coach three stores; no company-wide training

STILL TO CONFIRM · does the pattern hold in the four stores that already meet the target?

A request is a solution in disguise

A manager comes to L&D and asks for training. The sentence sounds like a problem statement. It is a solution, chosen before anyone established what is actually stopping the result. Reaching for training before finding the cause is the most expensive habit in the field, and the Dutch literature gave it a name: the costly misunderstanding (Arets & Overduin, 2008; Arets & Heijnen, 2008).

The request is easy to act on, which is precisely the trouble. It arrives with the solution attached, a number of participants, sometimes a date. A department can process it without ever asking which result is missing, in whose work, and how anyone would see it. The programme runs, the dashboard looks healthy, and the problem sits where it sat.

What makes the habit expensive is not the invoice. It is that the request names the one cause a course can address, and that cause is rarely the one holding the result back. Meanwhile the working hours disappear, in numbers nobody put on the request.

This paper describes the alternative, which is old, well documented and until recently almost impossible to run on every request: find the cause first, then decide. And it describes what changes when every learning advisor has a performance consultant beside them at the intake, instead of one specialist for the few requests there is time for.

The request says training. The cause almost never does.

What a request tells you

The number of participants, the format, sometimes a date. All of it about the solution. Nothing about the result that is missing.

What a diagnosis tells you

The business result the organisation wants, the gap in numbers, and the causes behind that gap, separated from the symptoms. That is what decides whether training is the answer at all.

Why the diagnosis loses

It costs time nobody budgeted, and it questions a manager's own reading of the problem. The quick answer suits everyone in the room.

What a request really costs



Illustrative figures. Direct cost per participant €300 to €500 for classroom teaching or €25 for e-learning; travel €120 to €200; working time at a loaded €40 an hour. The proportions carry the argument, not the totals.

Two things stand out the moment the numbers sit on the table. The invoice is the small number, and the large one is working time, which almost never appears on the request. And time scales without mercy. The e-learning for twenty thousand people looks cheap because it is digital; it consumes roughly 160,000 working hours.

None of this is controversial. It is arithmetic. What has changed is that the arithmetic now takes less time than reading the request it belongs to.

Not an argument against training

A department that sees these numbers and approves the training anyway is free to do so. What it can no longer say is that it did not know. That is the point of putting the number on the table before the decision instead of after it.

Where the cause really is a skill deficit, well designed learning is exactly right and nothing else will do. The discipline is diagnostic, not ideological.

Most of the leverage sits around the person

Thomas Gilbert built the discipline on one observation: a person's environment engineers performance far more than what sits inside their head (Gilbert, 1978). Knowledge, skill and motive matter, and they operate inside a work system that either makes good performance easy or quietly makes it impossible.

Four decades of quantitative work point the same way. A synthesis of meta-analyses across sixteen determinants of behaviour finds that knowledge, general skills, attitudes and beliefs have negligible effects as targets of an intervention; habits and social support have medium effects; and access, the material and logistic means that make a behaviour feasible, has the largest (Albarracín et al., 2024). A training request aims at the weakest determinant on that list.

One claim deserves a label. Practitioners in this tradition often say that only a minority of performance problems turn out to be genuine deficits of knowledge or skill. That proportion is a practice heuristic from fifty years of case work, not a meta-analytic finding [not verified]. The direction it points is well supported; the exact share is not.

the work around the person

expectations and feedback

tools, information, time

incentives and consequences

the person

knowledge and skill

capacity

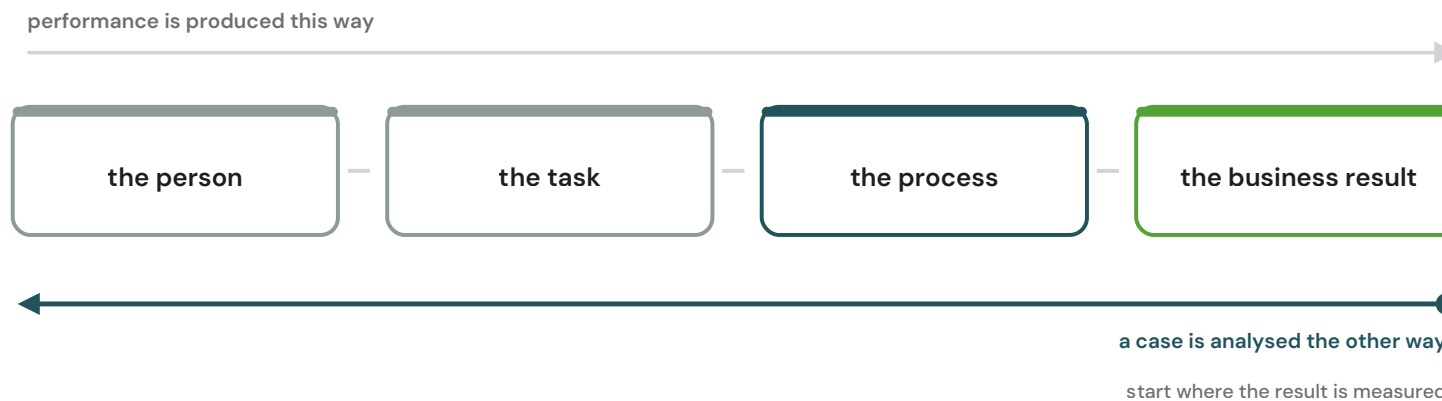
motives

the only box training can fill

The consultant looks in this order: first at the conditions around the work, which are the likeliest cause and the cheapest to repair, and only then at the person. Knowledge and skill is the single box a course can fill.

Work the case backwards

Performance is produced from left to right: people carry out tasks, tasks combine into processes, processes deliver results. A case gets analysed the other way. Start with the business result the organisation wants, then work back through the processes and the tasks to the place where performance actually leaks. Starting at the tasks, which is where the training request points, means starting at the wrong end (Rummler & Brache, 1995).



01 · The result

Stated as a measurable outcome, not as an activity. "More customer focus" is a wish. "Repeat orders at 62 percent against a needed 75" is a case.

02 · The gap

The quantified distance between the result now and the result needed, with a value attached. Only a number tells you whether the problem is worth solving and whether it got solved.

03 · The causes

Separated cleanly from the symptoms, drawn from more than one source, and never accepted on the strength of a confident opinion (Mager & Pipe, 1997).

None of this takes months. A good performance analysis is lean and fast, aimed at the few causes that matter rather than at an exhaustive audit. Its output is a short, evidence-based diagnosis that decides whether training is the answer at all.

Why so few functions work this way

If the case is this strong and this old, the obvious question is why almost no L&D function runs on it. Not because the method is in doubt. Because the daily pressures of the job run against it at every turn.

The field has known how to do this for decades. It has not been able to do it on every request.

No access to the real client

Between L&D and the manager who owns the problem sit intermediaries: HR business partners, competency owners, whoever handles training. The person who could describe the business priority is rarely in the room.

The confidence to challenge

Investigating a cause means questioning a manager's own diagnosis. Many advisors have neither the mandate nor the practice for that conversation.

The conspiracy of convenience

Everyone prefers the quick answer. The manager sees action, L&D gets a clear deliverable, and nobody has to sit with the possibility that the problem is the system they run.

Slow work in a fast culture

Diagnosis is careful, evidence-led thinking inside an organisation built to answer now (Kahneman, 2011). Advisors who have the skill get worn down by the current.

None of these barriers is an argument against the method. Each one is an argument for making the method cheap enough to run on every request, which is exactly what changed.

What the Performance Consultant AI does

The tool guides a conversation, not a form. It sits beside the learning advisor at the intake and runs the same five moves every time.

01 · The request

In plain language

It starts with one field and one sentence: “a two-day sales training for all ten thousand reps, is training the fix?” No intake form, no template to fill in first.

02 · The questions back

About numbers, not wishes

The tool asks where the problem shows up in the figures, who is missing which target, and where the revenue leaks. The advisor supplies the detail the organisation already has.

03 · The data

What the organisation knows

Documents, dashboards and loose figures go in. The diagnosis rests on the organisation’s own evidence rather than on an assumption about it.

04 · The gap, quantified

Cause by cause

Current and needed results side by side, and under each one the cause, tagged by factor, by level, and by how far the evidence carries.

05 · The advice

And the documents to use it

Four versions of the same diagnosis, edited by the advisor, ready for the meeting where the decision gets made.

The method does not change. What changes is the scarcity: an analysis that took a specialist weeks now fits inside the intake conversation, on every request rather than the few there was time for.

The anatomy of a diagnosis

Diagnosis

retail · 1,000 people · request: customer-experience training

DIAGNOSTIC CONFIDENCE
moderate to high

THE GAP

sales per visit €18.40, target €21.00, worth about €2.6M a year

WHAT IS CAUSING IT

■ out-of-stocks on the top 40 lines

process · triangulated

■ pricing errors at the register

tools · substantiated

■ no feedback on visit value per store

expectations · hypothesis

RECOMMENDED NEXT STEPS

fix replenishment and the price file first; coach three stores; no company-wide training

STILL TO CONFIRM · does the pattern hold in the four stores that already meet the target?

Illustrative diagnosis. The same anatomy whether the request covers fifty people or twenty thousand.

Every cause carries its confidence

Hypothesis, substantiated or triangulated. The tool marks its own uncertainty instead of hiding it behind fluent language.

It names what would settle the matter


The questions it would ask next, so that a strong diagnosis becomes a watertight one before anyone spends money.

It is advice, not a verdict

The diagnosis supports a decision that people make. It carries no guarantee of results, and the tool says so in plain words.


One conversation, four documents

The person who owns the money, the manager who owns the team and the analyst who owns the detail do not read the same document. The same diagnosis therefore comes out four ways, from one conversation, without a second analysis.



This is what changes the report L&D brings to the table. For years the function reported what it did: completions, participants, satisfaction. It can now report what it protected and what it made possible, in the language the board already speaks.

The dividend as a standing metric



Illustrative dashboard. The dividend reported in the same cadence as completions, skills coverage and engagement.

Every request a department redirects frees money and working time. Counted once, that settles an argument. Counted at every intake and rolled up, it becomes a standing monthly number per unit and per country, next to the metrics L&D already reports.


A saving is a claim, so it needs a rule. Count only what the organisation had genuinely committed, keep freed working time separate from money, label every estimate as an estimate, and never book the same saving twice. A number nobody can defend costs more credibility than it buys.

What counts as a dividend

The money and the working time an organisation does not spend on a solution that would not have moved the result, plus what it redirects to one that does. Nothing else belongs in the number.

Performance consulting is the crossing

The Arets L&D Business Models place four ways of working on one rising line. The two on the left begin after someone else has decided that learning is the answer. The Performance Partner begins before that decision, establishes the cause and prescribes what the business needs, training or not. The Performance Consultant AI is what makes that crossing survivable on a normal working day.



01

Order Taker

Takes requests and delivers courses. Attendance counts as success. The cause of the problem never gets established, so the result rarely moves.

02

Learning Specialist

Designs excellent evidence-informed learning, and starts only once somebody else has decided that learning is the answer.

03

Performance Partner

Establishes the cause first and prescribes accordingly: process, tooling, guidance, or training where training is genuinely the answer. Includes O2.

04

Value Creator

Performance consulting and evidence-informed organisational learning, reported in the organisation's own numbers. Includes O2 and O3.

The programme, the tool, the author

The programme

Tulser delivers performance consulting as one work, learn and implement programme for an L&D team: ten real requests from your own department, asked again; an advice per request that names what is going on, what helps and what it delivers; and one way of working the whole team uses afterwards. tulser.com

The Performance Consultant AI

The tool runs beside the intake conversation and is free to try. It replaces neither the L&D team nor the consultant: it removes the reason that used to decide the matter by default, which was never having the time for an analysis. tulser.ai

Jos Arets

L&D thought leader, author and founder of Tulser. He works at the intersection of corporate learning, performance improvement and evidence-informed innovation, in the Netherlands and internationally. j.aret@tulser.com

Selected sources

Albarracín, D., Fayaz-Farkhad, B., & Granados Samayoa, J. A. (2024). Determinants of behaviour and their efficacy as targets of behavioural change interventions. *Nature Reviews Psychology*, 3(6), 377–392. <https://doi.org/10.1038/s44159-024-00305-0>

Arets, J. (2026). Learning that crosses the levels: The Learning Conversion Framework [White paper]. Tulser. [DOI to be added]

Arets, J., & Heijnen, V. (2008). Kostbaar misverstand: van training naar business improvement. Academic Service. [no DOI available]

Arets, J., & Overduin, B. (2008). Liever (g)een training: op weg naar performanceverbetering. Academic Service. [no DOI available]

Binder, C. (1998). The six boxes: A descendant of Gilbert's behavior engineering model. *Performance Improvement*, 37(6), 48–52. <https://doi.org/10.1002/pfi.4140370612>

Gilbert, T. F. (1978). *Human competence: Engineering worthy performance*. McGraw-Hill. [no DOI available]

Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Straus and Giroux. [no DOI available]

Mager, R. F., & Pipe, P. (1997). *Analyzing performance problems: Or, you really oughta wanna* (3rd ed.). Center for Effective Performance. [no DOI available]

Rummler, G. A., & Brache, A. P. (1995). *Improving performance: How to manage the white space on the organization chart* (2nd ed.). Jossey-Bass. [no DOI available]

Stolovitch, H. D., & Keeps, E. J. (Eds.). (1999). *Handbook of human performance technology* (2nd ed.). Jossey-Bass/Pfeiffer. [no DOI available]

Cite as Arets, J. (2026). The request is not the problem: The Performance Consultant AI [White paper]. Tulser. DOI [to be added]

Also from Tulser: Learning that crosses the levels (the Learning Conversion Framework) and The four Arets L&D Business Models.

© 2026 Jos Arets and Tulser. All rights reserved. You may share this document in unchanged and complete form, with attribution. No adaptation, translation or reuse of parts without written permission.