TULSER

No number · © 2026 Tulser

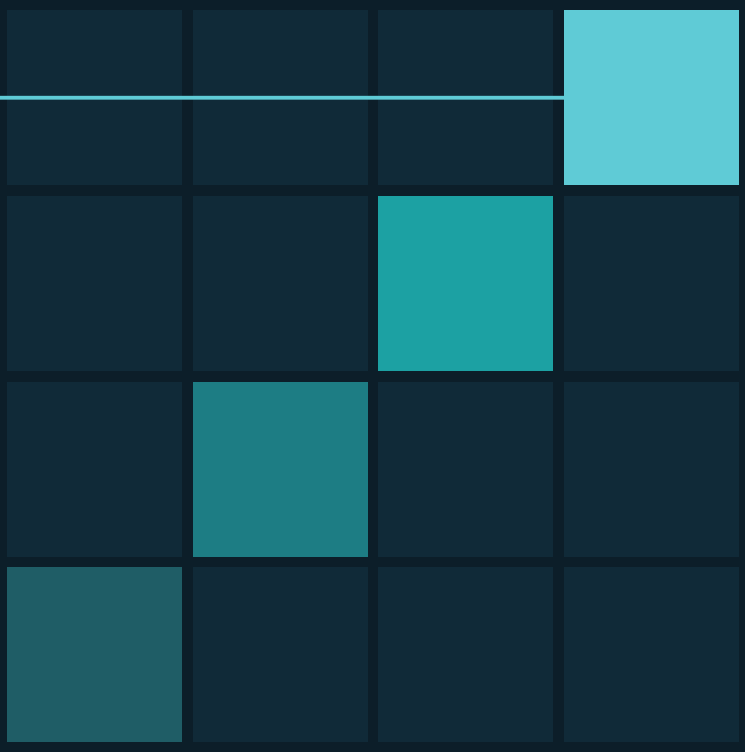
WHITE PAPERS · PERFORMANCE CONSULTING

TRAINING REQUEST · THE DESTINATION

Nobody in L&D wants to be an order taker. Now nobody has to be.

Arets L&D *Business Models*

The destination piece of the series, for the people who receive the requests. Four business models a learning function can run (Order Taker, Learning Specialist, Performance Partner, Value Creator) one divide between HR-driven and business-driven, and the route through the sixteen numbers from the left of that line to the right.



JOS ARETS

with Vivian Heijnen

IN THIS PIECE

Which model are you running?
And which do you intend to run?

THE QUESTION	*Every L&D function runs a business model, named or not. Which of the four is yours, which one do you intend to run, and which of the sixteen numbers carry you from one to the next?*
FOR	L&D leadership, learning advisors and HR business partners: the people who receive the training requests. Managers with a business problem do not need this piece; they enter at number 01.
READ	12 minutes · with a five-line instrument and a route map into the series.

01	The divide	Four models, one line: HR-driven or business-driven.
02	Which model are you running	How each model looks from inside a normal week.
03	The route through the sixteen	Which numbers carry each move, and where the instrument comes in.
04	Why the move is suddenly simple	Performance consulting held by an assistant, and the unnecessary training that no longer has to be delivered.
05	What this piece is not	Three boundaries, so the ambition stays honest.
.	Run it yourself	Five lines, on your own function.

Nobody in L&D wants to be an order taker.
Now nobody has to be.

Four models,
one line between them

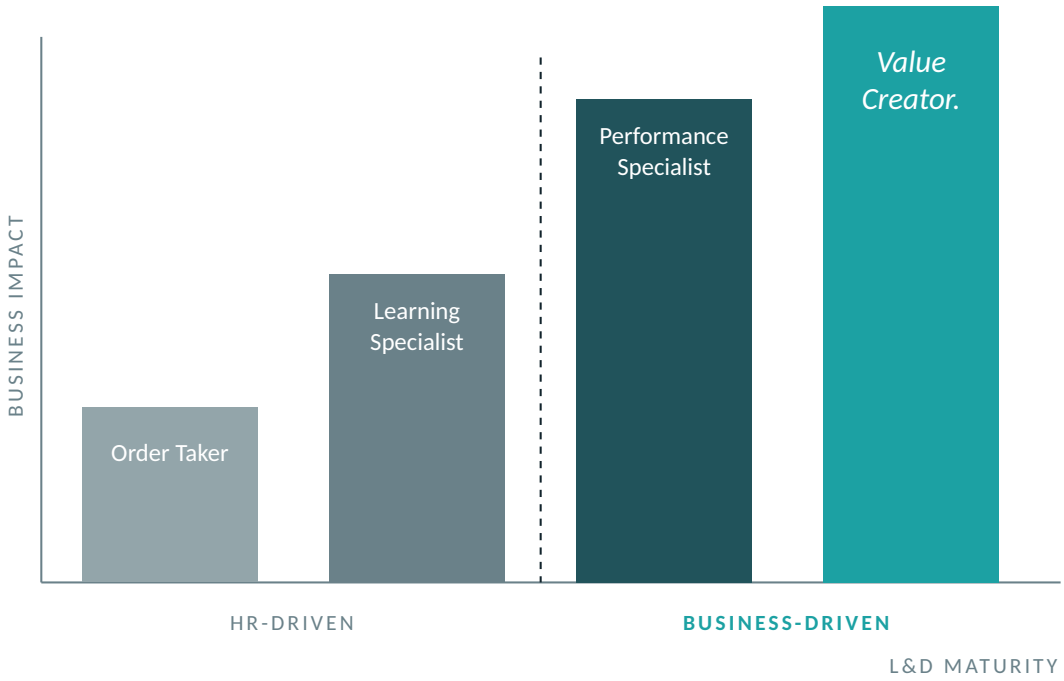
Every learning function runs a business model, whether anyone has named it or not. The Arets L&D Business Models name four, and they turn on one divide. On the left sit the HR-driven models, which develop individuals and report in learning metrics. On the right sit the business-driven models, which convert learning at the level of teams and the organization and report in the business metrics the organization already runs on (Arets, 2026).

Order Taker receives a request and fulfills it: off-the-shelf, eminence-informed learning as a service (bought on reputation rather than evidence, delivered on schedule, judged on completion). **Learning Specialist** is the professional standard for individual learning: evidence-informed, built to method, judged on the quality of the learning itself. Almost every learning function needs this model: it educates people in the field the organization works in, and it carries the license to operate that people must hold before they start. That makes individual learning the substrate: necessary, and the beginning of learning in an organization rather than the end. **Performance Partner** is the first business-driven model: it diagnoses the incoming training request before designing anything, serves team-based learning, and reports in the team's own business measure (model two's individual craft stays included). This is the level where Performance Consultant AI does its work. **Value Creator** runs the same discipline at the scope of the whole organization: organizational learning, with models two and three folded in, and the same instrument at hand. From the Learning Specialist upward the models are cumulative; the Order Taker stands outside that accumulation: it is the model the other three replace. The full model is set out in the Arets L&D Business Models white paper (Arets, 2026).

The divide is where the ambition lives. On the left, a function can be excellent and still unable to say, in the board's own language, what it returns. On the right, real business problems come onto the function's desk, and the conversation with the business happens, for the first time, in a single currency. Nobody in L&D wants to be an order taker. The route to the right used to be the hard part. That is what has changed.

The model, *drawn*

Four models, rising in maturity and in business impact from left to right, and one dashed line that carries the whole argument.



The Arets L&D Business Models. The vertical axis shows business impact, the horizontal axis L&D maturity. From the Learning Specialist upward the models are cumulative: the Performance Partner keeps model two's craft, the Value Creator keeps both; the Order Taker is the exception, carried along by none of them. The dashed line is the divide this piece is about: to its left the function develops individuals, to its right it converts learning where the business counts its results.

Which model
are you running?

Not which model the strategy document names, which one the weeks actually look like. Read the middle column and be honest; the right column is what the board hears.

MODEL	WHAT A NORMAL WEEK LOOKS LIKE	WHAT THE BOARD HEARS
Order Taker HR-driven	Requests arrive with the solution already named, and the shelf supplies it: learning as a service, eminence-informed rather than evidence-informed. The work is scoping, scheduling and sourcing. Nobody asks which result the course should move.	Completions, satisfaction, spend. Activity.
Learning Specialist HR-driven	Programs are designed to method, practice is supervised, transfer is measured: evidence-informed individual learning, educating people in the organization's own field and carrying its license to operate. The craft is real and necessary. The requests still arrive with the solution named.	The quality of the learning itself.
Performance Partner business-driven	Every request is treated as a hypothesis. A diagnosis of business, performance and cause precedes any design, and some answers are not programs at all.	The team's own business measure, moving.
Value Creator business-driven	The same discipline, organization-wide: what one unit learns is found, kept and reused by the next. The function sits in the conversation where results are chosen.	Organizational measures, and what did not have to be bought.

Four models, seen from inside a working week. The ranking is maturity, not a verdict on craft, where a function should land depends on what it is for, a choice, made and named.

The first strategic choice an L&D function makes
is about itself.

The route
through the sixteen

The sixteen numbers of this series are the route from the left of the divide to the right. Not all at once, and not in reading order: by move.

THE MOVE	WHAT CHANGES	WHERE THE SERIES CARRIES IT
Order Taker → Learning Specialist	The craft of individual learning becomes evidence-informed. The requests do not change yet; the quality of the answer does.	Number 11: learning metrics, and what they can honestly claim.
Learning Specialist → Performance Partner	The decisive move, across the divide. Every request becomes a hypothesis; diagnosis precedes design; the answer is chosen against a cause.	Numbers 12 and 05 to 10: the request, and the analysis behind it. Number 13 prices what the diagnosis prevents.
Performance Partner → Value Creator	Scope widens from the team to the organization: findings are kept, transferred and reused instead of rediscovered.	Number 11 for the currencies, number 14 for implementation at scale, and the memory the assistant carries between analyses.

Three moves, one route. The second is the one most functions never made, not for lack of will, but because every diagnosis cost a week and a practitioner.

WHAT MAKES THE ROUTE PASSABLE NOW

The move across the divide stands or falls with one discipline: a diagnosis before every design. That used to be unaffordable at the pace requests arrive, which is why the left side of the line is so crowded. **Performance Consultant AI makes the diagnosis fast, structured and repeatable**: run on the request in front of you, in minutes once the data is in. Number 15 shows the instrument at work. The instrument makes the route passable; choosing the destination remains yours.

Performance consulting, *held by an assistant*

Performance consulting was never conceptually hard. It was operationally out of reach: five questions in a fixed order, records to pull and reconcile, a gap to compute and cut, an exemplary performer to find (a week of work per request, at the pace of a function that receives requests every day).

That is the part the assistant now holds. Performance Consultant AI carries the five questions in order and will not run past an unanswered one, pulls and reconciles the figures, cuts the gap, sets the request beside the data nobody had joined up, and drafts the reports in the business's own currency, with a confidence rating that says how sure it is, and what would make it surer. The judgement stays with the advisor: which measure is right, which cause is binding, whether to sign. What disappears is the week that made the judgement unaffordable.

And that changes what the function produces. The request that cannot work no longer becomes a program; it becomes an evidence-backed no, with a better answer attached: a view, a rota, a card, a process fix. The request that is justified gets a sharper yes, aimed at a diagnosed cause. **The unnecessary training simply stops being made**, and every euro and every working hour not spent on it is capacity the function keeps: for evidence-informed individual learning where that is the answer, and for the team- and organization-level work that the business-driven models exist to do. That is why the instrument strengthens L&D in the conversion from HR-driven to business-driven: it does not shrink the function, it frees the function to do the work that pays.

What the board then receives is not a course catalogue but a decision document: advice, the bottom line, the money, and how sure the diagnosis is (number 15 shows the instrument in full).

What this piece *is not*

Three boundaries: the limits that keep the ambition honest.

- **It is not a verdict on where you stand.** The ranking is maturity, not craft. An organization that asks its learning function for one dependable supply of competent people is right to run a mature Learning Specialist, as a choice, made and named.
- **It is not the door for managers.** This piece exists for the people who receive training requests. A manager with a business problem does not need a destination for L&D; they need a diagnosis of their own result, and they enter the series at number 01.
- **It is not a promise that an instrument moves you.** Performance Consultant AI makes the diagnosis routine; the decision to answer to the business, and to be measured in its currency, is a strategic choice no tool can make for you.

RUN IT YOURSELF

Five lines,
on your own function

Fifteen minutes, honestly. Line five is the one that turns an answer into a strategy.

1 · The last five requests

Did they arrive with the solution named? How many were diagnosed before they were designed?

2 · The currency

Take your last report to management. Learning metrics or business metrics?

3 · The model, named

Which of the four does that make you today? Not the strategy document: the weeks.

4 · The intention

Which model does the organization actually need from you? Ask what it funds you for, not what it says.

5 · The first move

Take the next request that arrives and run the diagnosis before you answer it. Number 12 is the script; number 15 is the instrument.

WHERE THIS PIECE SITS

Three orders, *one series*

The series has three orders, and this piece is the first of them. ***The destination****: this piece, unnumbered, for the people who receive the requests: where do you want your function to be?* ***The question you are on****: the sixteen numbers, the method itself, one step at a time.* ***The form on your desk****: Solutions that cannot work, the field guide, last in the set and first in use: seventeen forms a performance problem arrives in, each with an address in one of the numbered papers. Managers and staff with a business problem skip this piece and enter at number 01; the destination is for the function that answers them.*

HOW THIS SERIES CAME ABOUT

Human performance technology travels the way a craft does: from hand to hand. Vivian Heijnen and Jos Arets learned the analysis from Dennis Mankin, in San Diego, more than twenty years ago, and have kept in contact with him ever since. Both remain grateful to him: for the rigor of the method as he taught it, and for the leadership he has given performance consulting over a long career.

What followed was twenty years of use. Vivian took the method into organizations in the Netherlands and well beyond them, and what accumulates over that many analyses is largely tacit: the cut nobody thought to make, the question that opens a defensive room, the finding that turns out to be an artifact of the data. Making that knowledge explicit, case by case, is what this series is for. Jos wrote it down as it went, in two books and more than forty articles and blog posts, most of them in Dutch, and one seven-part English series, *Is AI the bicycle of the mind?*, which number 15 grows out of. The books: *Liever (g)een training* with Boudewijn Overduin, *Kostbaar misverstand* with Vivian Heijnen, which this series takes a good deal of its thinking from.

None of that is reprinted here. It has been reworked: brought up to date, tested against the research that has appeared since, and dropped where the evidence no longer carries it. Theory that has never met a client is a position; practice that cannot account for itself is a story. This series is what remains when each is held to the other. Performance Consultant AI carries the same discipline to the desk where the decision is actually made.

REFERENCES

Where every claim *in this piece comes from*

Arets, J. (2026). *Arets L&D Business Models: A future-proof framework for corporate L&D* [White paper]. Tulser.

This piece announces; the white paper on the models (Arets, 2026) explains. The evidence rule of the series is set out in number 02, and number 14 carries implementation at scale.