

A TULSER WHITE PAPER

The strongest lever is access

Access to expertise at the moment of need: what the research says, how an organization delivers it, and where the boundary lies.

One argument, two delivery channels

An AI platform for a whole profession, and the Expertise Amplifier for one theme. Both aim at the same lever, both serve the individual, the team and the organization, and both leave the boundary intact.

THE EVIDENCE · THE LINEAGE · THE BOUNDARY ·
BEYOND THE INDIVIDUAL

Jos Arets · Tulser

White paper · September 2026



L&D buys knowledge where the work needs access

Ask an L&D function what it delivers and the answer is about knowledge: modules, learning paths, a course that brings people up to date. Ask a nurse, an engineer or an advisor what they need at the moment itself and the answer is about something else. The right rule, the right threshold, the agreement that changed last month. Not in advance. Now.

That difference is not a detail. It decides where the money goes. A function that thinks in knowledge buys courses and measures completions. A function that thinks in access designs the work so that what people need is within reach when the question arrives, and measures whether the result moved.

Neither is a matter of taste. The research on what changes behaviour has a clear ranking, and it is uncomfortable for anyone whose product is a course. Knowledge as the target of an intervention does close to nothing. Access does the most.

One more thing changes the arithmetic. Almost everyone in work already holds the licence to operate: the qualification, the diploma, the registration. After that, formal training is genuinely required where the law or the sector prescribes it, and that is the compliance route. Everything else is not a training question at all. It is a question about the cause of a performance problem, which an advisor establishes through performance consulting, with the Performance Consultant AI alongside (see The request is not the problem).

This paper puts four things side by side: what the evidence says about access, where the idea comes from, how an organization delivers it without a two-year project, and where the boundary lies. The boundary belongs in the argument, not in the small print.

Knowledge in advance is an assumption. Access at the moment of need is a design.

What a course assumes

That the right knowledge is present, remembered and retrievable weeks or months later, in a situation nobody could describe in advance.

What access assumes

That the question arrives in the work, and that the answer is one step away, in a canon somebody keeps current, so a whole team gets the same one.

Why this is a money question

Working time is the largest cost of any course and it never appears on the request. Access costs a licence and a security review.

Sixteen determinants, one clear ranking

A synthesis of meta-analyses across sixteen determinants of behaviour finds a ranking that is hard to argue with. Knowledge, general skills, attitudes and beliefs have negligible effects as targets of an intervention. Habits and social support have medium effects. Access, the material and logistic means that make a behaviour feasible, has the largest (Albarracín et al., 2024).


Read that against a training catalogue and the conclusion is blunt. Most of what L&D sells aims at the weakest determinant on the list, and does so at the highest cost, because the working hours it consumes never reach the invoice.

What is established here, and what is not. The ranking above comes from research on the behaviour of individuals, and it holds. That interdependent work also needs access and support at the level of the team and the organization was not measured in those studies. It is the inference I draw from them, and it is a proposition still to be tested (Arets, 2026).



The idea is older than the technology

Gery described the destination in 1991. What arrived late was the vehicle.



Performance support is not a new word, and moment of need is not a new idea. Gloria Gery gave both their shape in 1991: a system that puts the right knowledge, in the right amount, within reach of the person doing the work (Gery, 1991).

Two things in that definition matter now. It was information, and it was passive. Documents, procedures, step-by-step guidance, mostly for procedural tasks. And it leaned on technology that in 1991 barely existed, so the idea outran its execution for thirty years.

It had a purpose that is easy to forget. For Gilbert and for Rummler and Brache, support inside the work was the short way home: what you reach for when the cause was never a deficit of skill (Gilbert, 1978; Rummler & Brache, 1995).

What AI adds is not another channel but the integration. Looking something up, having it explained and being walked through the exception now happen in one place, and that place strengthens what the professional already knows how to do. That is augmentation rather than support (Arets, 2023a, 2023b).

Three models L&D runs on, and all three aim at one person



Five moments of need

Gottfredson and Mosher took Gery's idea and named the moments at which a person needs support: learning something for the first time, expanding it, applying it, solving a problem, adapting to change (2011). Useful vocabulary for when support is needed.

70:20:10

Moved the attention from the course to the work, and that was a necessary correction (Arets, Jennings & Heijnen, 2015). It still counts what one person learns, wherever that learning happens.

ADDIE


Gives design a procedure from analysis to evaluation, and it still organizes the work of most L&D departments (Branch, 2009). The object of that design is one intervention and its participants.

These three are what the field runs on, and they share one limitation, which is structural. Each describes one person meeting one task. None has a place for the team that has to reach one shared answer, or for the organization that has to keep what it works out, and the environment around the work sits outside all three.

This is a position, not a finding. All three are serviceable within their own scope, and 70:20:10 is work of Charles Jennings, Vivian Heijnen and mine, so the criticism includes it. The claim is that a model of learning has to account for the levels above the individual as well, and that none of these three does.

The question does not arrive in the classroom

A course sits weeks before the work. The question arrives inside it: at the bedside, at the machine, in the intake conversation. Everything between those two points is memory, and memory is exactly the part the evidence says not to lean on. Access moves the answer to where the question is.



At the bedside

May a care worker administer this on their own?
The answer sits in a working agreement that changed in June. Nobody knows that by heart, and nobody should have to.

At the machine

The step after the exception, the torque for this variant, the form that goes with it. Trained once, needed twice a year.


At the intake

Which question to ask a manager who wants a training. That is expertise, not a document, and it is exactly what an assistant can hold.

What is learned and not used, fades

Skills decay, and the loss grows with the time a skill goes unused (Arthur et al., 1998). That is not a criticism of learning. It is a statement about where learning belongs in time. Everything that only comes up twice a year is a poor candidate for a course and a good candidate for access.

Organizations lose knowledge in the same way. What one team works out disappears when the roster changes, unless somebody stores it where it can be found and keeps it current (Argote, 2024). An assistant that answers from a governed canon is one of the few places where that maintenance is somebody's job by design.



What may be looked up, and what may not

The argument only holds with its boundary attached. Some things have to be there without looking: acting under time pressure, reasoning your way through an unfamiliar fault or an unfamiliar patient, coordinating with a team, seeing that a case is unusual before anyone says so. Those need whole-task practice, deliberately designed and built up over time (Van Merriënboer, Kirschner & Frèrejean, 2024). No assistant produces them.

what may be looked up

- a rule that changed last month
- a threshold, a dose, a torque, a deadline
- the step after the exception
- which form, which route, which owner

what has to be there without looking


- acting under time pressure
- clinical or technical reasoning
- coordinating with a team
- seeing that a case is unusual

The line between the two is a design decision, taken per task, and it is the first question in any serious conversation about AI in learning.

Two more things stay outside. An assistant produces no record of demonstrated competence, so proving competence remains its own route. And an answer is worth exactly the canon it rests on: an assistant on an ungoverned pile of documents replaces a knowledge problem with a trust problem.

The step after e-learning and performance support

Nothing here breaks with what the profession already does. E-learning solved reach. Performance support moved the answer into the work (Gery, 1991). What changes is what an assistant does with a governed canon: it answers in the language of the question, stays current because one editorial board keeps the canon, and serves the team.



Good digital learning exists

Design whole tasks the way Ten Steps to Complex Learning prescribes, and follow the multimedia principles (Van Merriënboer, Kirschner & Frèrejean, 2024; Mayer, 2020). That practice is the other half of the boundary, not a competitor.

What changes, and what does not

The aim is the one Gery wrote down: the right knowledge, in the right amount, at the moment of need. Three things differ. The answer arrives in ordinary language, it stays current because one editorial board keeps the canon, and one canon serves individual, team and organization.

How to choose, per task

Ask two questions. Must someone do this without looking it up? Then design practice. Does it change often, or come up twice a year? Then design access. Tulser delivers both, so the answer does not depend on what we happen to sell.


Support for one person is where this starts, not where it ends

Classical performance support helps one person finish one task. That is worth having, and in interdependent work it is not enough. A shift rarely fails because one technician cannot find an instruction. It fails because two people hold different versions of the same agreement and neither of them knows it.

A governed canon works at three levels at once. The individual gets the answer at the moment of need. The team gets one answer instead of three, which is what makes coordination possible in the first place. And the organization keeps what one team worked out, because the answer lives in a canon somebody owns and updates rather than in the memory of whoever was on shift.


That third level is where access meets the Learning Conversion Framework. What a team learns has to outlive the team, and a canon somebody owns and updates is one of the few forms in which it reliably does (Arets, 2026).

The character of the support has changed too. Gery described passive information, retrieved by someone who already knew what to look for. A governed canon inside an assistant answers the question in the words of the person asking it, and gives the reasoning with it, so the judgement stays with the professional. That is augmentation rather than a manual.



Two channels, one assistant, one boundary

Access is not a product, it is a property of the work. Two routes deliver it, and the choice between them is technical rather than philosophical. The model behind both is old enough to have been studied: medicine has run a governed, licensed canon for thirty years, and a study across US hospitals finds a small but consistent association with shorter stays and lower risk-adjusted mortality (Isaac, Zheng & Jha, 2012). An association, not a cause.



An AI platform

A composed body of knowledge for a whole profession, extended where needed with the organization's own material. Certification, security, management and licences come with it, so it is a project of months, delivered by an IT partner.

An Expertise Amplifier

One theme, delivered as a read-only connection to an assistant the organization has already approved. Light on IT, delivered in weeks, priced as a licence per year.

How to choose

A profession that works with protocols every day earns the platform. One theme, or a fast start, or an assistant that already stands, points to the Amplifier. They do not exclude each other.

Four steps, and the first pays for itself

01 · Sources

What applies and what has lapsed

Almost every organization finds contradictions and expired documents in this step. That gain arrives with or without AI, which is why it is the first step and not the last.

02 · Structure

One file per theme, ordered by the work

What survives step one rarely amounts to a document per question, so the editorial board writes it: one plain file per theme, linked to the files beside it, ordered on questions, tasks and moments rather than on the folder a document happened to sit in. Here an assistant is usable rather than merely available.

03 · Stewardship

Who writes, who checks, when

Without that agreement a canon goes stale within a year and the assistant quietly becomes unreliable. Stewardship is a job, not a phase.

04 · Adoption

Start with the critical tasks

Begin where it matters most: the tasks that go wrong often, the ones whose consequences are large, the questions people ask every shift. Load the expert knowledge those tasks need, available to everyone around the clock, and keep it current together with the organization. Use follows from usefulness; a launch communication has never made an assistant useful.

An organization that runs these four steps ends up with two things that last: a canon that holds and an editorial board that keeps it current. You replace a platform; you keep a canon.

One editorial board

Somebody writes while everyone else reads, and a log says what changed and when.

Read-only

The connection reads and never writes. What a team records about its own work stays with the team.

No personal data

Nothing about individuals goes into the canon, which keeps the security review short.

A brake on extraction

Limits in the connection and in the contract. An assistant asked to print the whole canon is a real risk.

What this changes on Monday

Ask where the question arrives

For every request: does this knowledge have to be in someone's head, or within reach at the moment of need? The answer decides the intervention.

Split the catalogue in two

What may be looked up leaves the course and becomes access. What has to be there without looking gets real practice, and more of it.

Design above the individual

Ask what the team needs to share and what the organization has to keep. That question is what separates access from a help file.

Own the canon, not the platform

The lasting asset is the governed knowledge and the editorial discipline behind it. Measure what got decided differently, not how many questions were asked.

The two channels, the framework, the author

The two channels

AI platforms: a composed body of knowledge for a profession, set up by Tulser and delivered by an IT partner. The Tulser Expertise Amplifier: one theme, delivered by Tulser as an MCP server on the assistant you already use. tulser.com

The framework underneath

Where access meets team and organizational learning, the Learning Conversion Framework supplies the mechanism: Learning that crosses the levels (Arets, 2026). The four Arets L&D Business Models place the same move on one rising line. Both free at tulser.com.

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.aretst@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. (2023, September 7). Is AI the bicycle of the mind? Performance support powered by artificial intelligence. *eLearning Industry*. <https://elearningindustry.com/is-ai-the-bicycle-of-the-mind-performance-support-powered-by-artificial-intelligence>

Arets, J. (2023, September 13). Is AI the bicycle of the mind? Performance augmentation in corporate eLearning. *eLearning Industry*. <https://elearningindustry.com/is-ai-the-bicycle-of-the-mind-performance-augmentation-in-corporate-elearning>

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

Arets, J., Jennings, C., & Heijnen, V. (2015). 702010 towards 100% performance. *Sutler Media*. [no DOI available]

Argote, L. (2024). Knowledge transfer within organizations: Mechanisms, motivation, and consideration. *Annual Review of Psychology*, 75, 405–431. <https://doi.org/10.1146/annurev-psych-022123-105424>

Arthur, W., Jr., Bennett, W., Jr., Stanush, P. L., & McNelly, T. L. (1998). Factors that influence skill decay and retention: A quantitative review and analysis. *Human Performance*, 11(1), 57–101. https://doi.org/10.1207/s15327043hup1101_3

Branch, R. M. (2009). *Instructional design: The ADDIE approach*. Springer. <https://doi.org/10.1007/978-0-387-09506-6>

Gery, G. J. (1991). *Electronic performance support systems: How and why to remake the workplace through the strategic application of technology*. Weingarten Publications. [no DOI available]

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

Gottfredson, C., & Mosher, B. (2011). *Innovative performance support: Strategies and practices for learning in the workflow*. McGraw-Hill. [no DOI available]

Isaac, T., Zheng, J., & Jha, A. (2012). Use of UpToDate and outcomes in US hospitals. *Journal of Hospital Medicine*, 7(2), 85–90. <https://doi.org/10.1002/jhm.944>

Mayer, R. E. (2020). *Multimedia learning* (3rd ed.). Cambridge University Press. <https://doi.org/10.1017/9781316941355>

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]

Van Merriënboer, J. J. G., Kirschner, P. A., & Frèrejean, J. (2024). *Ten steps to complex learning: A systematic approach to four-component instructional design* (4th ed.). Routledge. [no DOI available]

Cite as Arets, J. (2026). The strongest lever is access [White paper]. Tulser. DOI [to be added]

Also from Tulser: Learning that crosses the levels, The request is not the problem 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.