TULSER

No. 08 · © 2026 Tulser

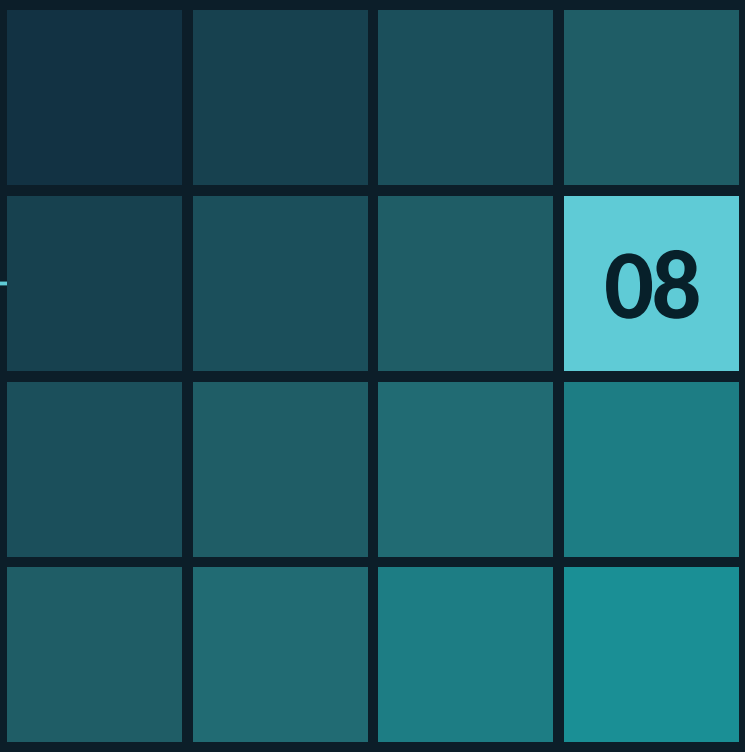
WHITE PAPERS · PERFORMANCE CONSULTING

BUSINESS PROBLEM · QUESTION 4

The gap is established. Now: at which level does its cause actually sit?

Cause analysis: *the three levels*

Gilbert's unit of analysis and the chain of accomplishments it produces, the three levels that chain runs through, and the white space between them, because a remedy aimed at the wrong level cannot work, however good it is.



JOS ARETS

with Vivian Heijnen

IN THIS PAPER

Whose accomplishment *is actually short?*

THE QUESTION — *The gap is measured and somebody already achieves the level. At which accomplishment in the chain does the shortfall first appear, and does its cause sit in the organization, in the process that crosses it, or in the job as one person does it?*

FOR — Anyone who owns a result or is asked to fix one: line managers and directors, operations, quality and continuous improvement, process analysts, OD, HR business partners and L&D.

READ — 24 minutes · with a worked case and a one-page instrument.

01	Whose accomplishment, and what does it feed?	Behavior or accomplishment, and why the choice decides everything after it.
02	Three levels, and why the word is not a metaphor	A cause at one level cannot be reached by an intervention at another.
03	The level nobody owns	The white space: what the evidence does and does not say about it.
04	The process as drawn and the process as run	Work-as-imagined, work-as-done, and why the difference is not disobedience.
05	The instrument that changed	Process mining: what it makes visible, and what it cannot tell you.
06	Four tests that locate a level	None of them takes longer than a day, and the first one is free.
07	The insurer whose handlers were not slow	A worked case: every person met the standard, the result did not.
08	What this does not claim	Four limits, because the rest is not credible without them.
.	Run it yourself	Six lines, on the gap in front of you.

The organization chart shows the boxes.
The result is produced between them.

Whose accomplishment, *and what does it feed?*

Before any scheme of levels, Gilbert's question, and it is worth setting his move out properly, because the rest of this series stands on it and it is almost always summarized into meaninglessness.

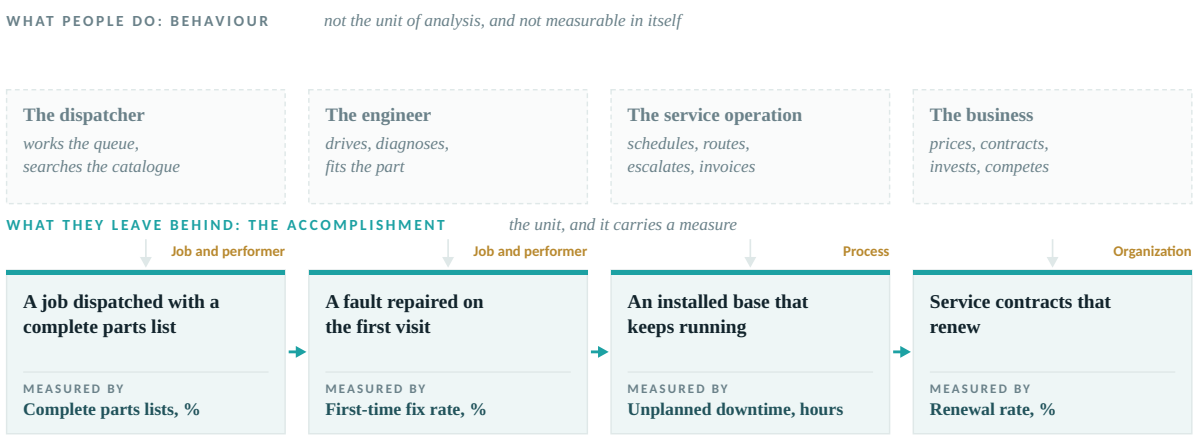
Gilbert built the discipline on a distinction between *behavior* and *accomplishment*. Behavior is what somebody does: the dispatcher works a queue and searches a catalogue, the engineer drives and diagnoses and fits a part. An accomplishment is what is left behind when they stop: a job dispatched with a complete parts list, a fault repaired on the first visit. Gilbert's claim is that the second is the unit of analysis and the first is not (Gilbert, 1978).

Three consequences follow, and each one does work later. An accomplishment carries a measure and a behavior does not: you can count complete parts lists, and "searching the catalogue properly" is not a quantity. An accomplishment is what the next person actually receives, so it is the unit that can be short. And two people can produce the same accomplishment by entirely different behavior, which is why copying what an exemplary performer does is the wrong instruction and reading what they produce is the right one: the subject of number 07.

> **THE TEST, IN ONE QUESTION**
>
> Ask of anything somebody proposes to improve: *if this went perfectly, what would be lying on the desk afterwards that is not there now?* If the answer is a noun with a number attached, you have an accomplishment. If the answer is a description of effort (better communication, more ownership, greater awareness) you have behavior, and there is nothing yet to analyze.

And accomplishments *nest*

That move settles the level question before it is asked, because accomplishments are not free-standing. Each one is somebody's output and somebody else's input, and they run in a chain to the number the board reads.



Four accomplishments, each of them somebody's output and the next one's input. Ask where the cause sits at the wrong link and every answer will be plausible and useless.

The top row is what a walk round the floor shows you. The bottom row is what can be measured, what can be short, and what the next link receives. Gilbert's whole argument is the arrow between the two rows.

Which link
is actually short?

So “where does the cause sit” is not a philosophical question about systems. It is a concrete one: **at which accomplishment in the chain does the shortfall first appear?**

Ask it at the wrong link and every answer will be plausible and useless. That is what happens when a renewal rate is discussed in a room of dispatchers, and equally when a dispatcher’s output is treated as a problem of strategy. Both conversations can run for an hour and produce sensible-sounding proposals, and neither can move anything, because the accomplishment under discussion is not the one that is short.

The chain also explains why a number can look healthy three links up while something is badly wrong below it. An installed base can keep running on overtime and second visits, so unplanned downtime holds while first-time fix falls, and the cost surfaces somewhere else entirely, in a quarter nobody has connected to it. Following the chain downwards until the measure stops being healthy is the fastest way to find the link that matters.

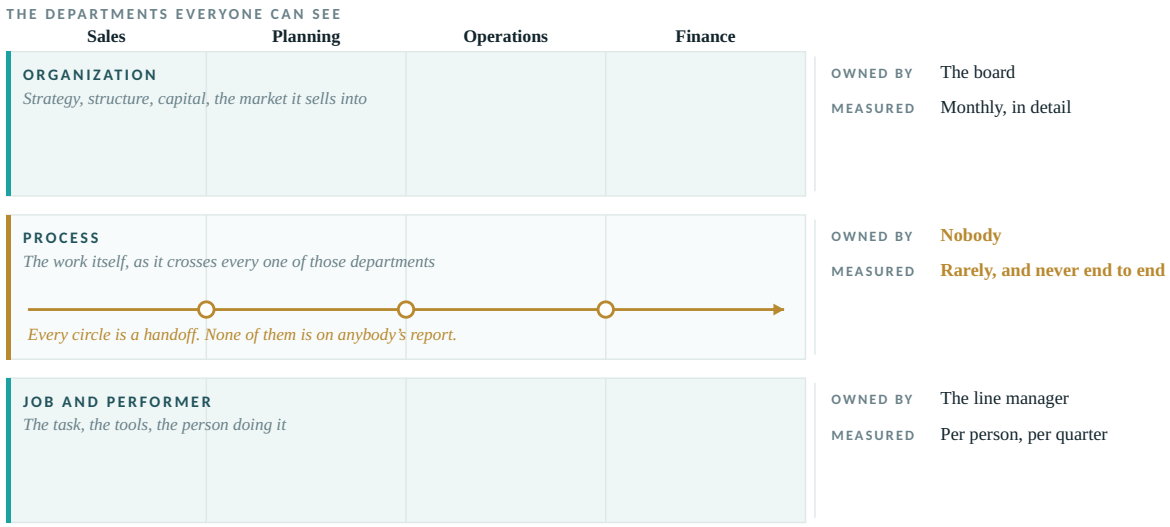
> **TWO MODELS, AND WHICH ONE THIS SERIES LEADS WITH**
>
> Gilbert’s question is *whose accomplishment, and what kind of cause is holding it back*: the six conditions in number 09. Rummler and Brache’s question is *where* that cause sits, at the organization, the process or the job. Both are sound and this series uses both. We lead with Gilbert, for two reasons: his unit is the accomplishment, which is what an organization already counts, and his six conditions are the categories an organization can actually change. Rummler’s levels earn their place as the sharper instrument for one specific job: finding the middle level, which the chain above shows but does not make an issue of.

Three levels,
and why the word is not a metaphor

Rummler and Brache put performance on three levels: the organization, the process that runs through it, and the job as one person does it. The distinction sounds like a management diagram. It is closer to a load-bearing wall.

The reason is that an intervention reaches exactly one level. Coaching a handler changes what one person does; it cannot change a queue that sits between two departments. Redesigning a handoff changes the route a file takes; it cannot change a target that was set too high in the first place. Every remedy has a reach, and the reach is a level. Choose the wrong one and the money is spent, the effort is real, and the number does not move, which is the outcome described in number 01, and it is almost always this mistake in another set of clothes.

The idea that a level is a formal property rather than a manner of speaking is not peculiar to this discipline. Organizations are nested: what happens above shapes what is possible below, and what happens below combines into something above that is not simply the sum of it. The consequence that matters here is a plain one: a claim established at one level does not automatically hold at another, and treating it as though it does is how a good finding produces a useless remedy.



The organization chart draws the columns. The result is produced by the row that crosses them, and that row has no owner and no report.

The columns are what the organization chart draws and what the reporting follows. The row that produces the result runs across all of them.

What each level *looks like from inside*

Three levels, three vocabularies. A problem at each one arrives in the room sounding entirely different, and the sound is a usable clue.

LEVEL	HOW THE PROBLEM IS DESCRIBED	WHAT CAN ACTUALLY REACH IT
Organization	"We are in the wrong segments." "The target was never achievable." "Two divisions are selling against each other"	A decision about strategy, structure, capital or capacity. Nothing below this level touches it, and no amount of local excellence compensates for it.
Process	"It takes forever and nobody can tell me why." "It goes back and forth between us and them." "By the time it reaches me the information is a week old"	A change to the route, the sequence, the rule that governs a handoff, or the information that travels with the work. This usually requires two departments to agree.
Job and performer	"He does not do the check." "She is much faster than the others." "They were never trained on the new system"	Anything that reaches one person at a time: the tool, the instruction, the feedback, the authority, the certification. Number 09 sorts these into six.

The vocabulary of the complaint is the first evidence about its level, and the third row is the one that arrives most often, whatever the level turns out to be. *[practitioner observation, not a measured frequency]*

> **WHY THIS QUESTION SITS FOURTH**
>
> The three before it did the work that makes this one answerable. Number 05 established which result is at stake and who reports it. Number 06 turned it into two values and cut the gap at least three ways. Number 07 found somebody already achieving it. Asked first, "where does the cause sit?" is a debate about opinions. Asked fourth, it is a question about a distribution somebody has already looked at.

One caution before the rest. Levels are not exclusive. A gap can have a cause at the process level and a second one in the job, and the honest finding names both with their shares rather than choosing the more convenient. What the question forbids is the reflex: assuming the lowest level because it is the one you can see, which is the subject of number 04.

The level *nobody owns*

Rummler and Brache called the space between the boxes on the chart the white space, and argued that most results are lost there. It is the most quoted claim in the discipline and the least tested, so it is worth saying carefully what the evidence supports.

Start with what is structurally true and needs no study. A department has a head, a budget, a report and a meeting. The interface between two departments has none of those. Nobody is appraised on it, nothing measures it end to end, and the work that fails there fails in a place with no owner. That is not a claim about frequency. It is a claim about attention, and it follows from the chart.

The measured evidence is thinner and points the same way. Pooling thirty years of research on how teams manage their boundaries, the activities by which a unit establishes and regulates its links with everything around it, a meta-analysis of 85 studies covering 10,848 teams found a relationship with team performance of $\rho = .39$, with outward-facing spanning activity doing better than protective buffering (Leicht-Deobald et al., 2025).

That number needs its limits stated in the same breath. It is about teams rather than about processes crossing four departments, it aggregates largely correlational studies, and its effect is strongest for boundaries with the outside world rather than for the internal handoff. It supports the direction of the argument and it does not prove it.

> **WHAT WOULD SETTLE THIS**
>
> There is no study that partitions performance problems across the organization, process and job levels and reports the shares. It could be done: take a few hundred completed analyses from consultancies and internal teams, classify the established cause by level against a fixed coding scheme, and report the distribution with the coding reliability. Until somebody does it, anybody quoting a percentage for “most problems are systemic” is quoting a conviction. This series argues the order of examination, which does not need the percentage, and says so.

Where the evidence *is actually strong*

On the interface itself the evidence is not thin at all, and the largest synthesis of it has nothing to do with hospitals.

Eighty independent samples drawn from eighty-six studies, covering 17,467 observations, were pooled to test whether integration across the boundary is related to how a firm performs: integration meaning the deliberate linking of processes, information and decisions, internally between functions and externally with customers and suppliers. The relationship is positive and significant (Leuschner et al., 2013). It is correlational, and it measures integration rather than redesign. But it is manufacturing, distribution and services rather than care, and it says plainly that the boundary is where a large part of firm performance is decided.

For what to *do* at that boundary, the closest look comes from the setting that has studied it longest, because there the consequences are visible within the hour. A systematic review of systematic reviews associates poor clinical handover with information omissions, errors and delays, and finds that the interventions that work are standardizations of the interface itself: a fixed form, a fixed sequence, a record that travels with the work (Desmedt et al., 2021).

Two consequences follow, and the second is the useful one. First, an interface can be redesigned and the redesign can be shown to work: the white space is not merely a diagnosis, it is an address you can send a remedy to. Second, note what the successful interventions were not. They were not training the people to hand over more carefully. They were changes to the artifact that crosses the boundary. That is a process-level remedy for a process-level cause, and it is the form of almost every fix at this level.

> **AND THE LITERATURE THAT CANNOT BE ADDED TO THIS**
>
> Business process redesign has a literature thirty years deep and it is almost entirely case reports and narrative reviews of methods, so under the evidence rule of this series none of it can carry a number. That is not a reason to distrust redesign; it is a reason to be exact about what is being leaned on here. Two syntheses and one clinical review, whose authors name the limit on generalising themselves, and the structural argument, which needs no study at all.

> *Without a Process Owner, the "white spaces" tend to be ignored.*

GEARY A. RUMMLER & ALAN P. BRACHE

Improving Performance, 3rd ed., 2013, p. 214

The process as drawn
and the process as run

Every organization has two processes with the same name. One is in the quality system, drawn as a flow with seven steps. The other is what actually happens, and nobody has a picture of it.

Safety research gave the distinction its working names, work-as-imagined and work-as-done, and they are worth borrowing. Work-as-imagined is the process as it is planned in advance or reconstructed afterwards, by people who are not doing it at the time. Work-as-done is what people actually do, under the conditions that actually obtain, including the conditions the design did not anticipate.

The gap between the two is not disobedience and it is usually not carelessness. It is what competent people do when the design does not fit the day: the extra spreadsheet that carries the information the system will not, the informal call that skips two steps because the formal route takes three days, the check performed at the end because performing it at the start would stop the line. And the resolution, where anybody attempts one, is almost never enforcement of the written version. It is redesigning the procedure with the people who were working around it, which is a process-level remedy arrived at by asking rather than by instructing.

> **WHY THIS MATTERS FOR LOCATING A LEVEL**
>
> An analysis conducted against work-as-imagined finds nothing, because on paper the process is fine. It then reports that the process is fine and moves down a level, where there are people to look at. That is the single most common route from a process-level cause to a job-level remedy, and it is why this paper insists on evidence of the process as run before any conclusion about the process at all.

Which raises the practical problem the discipline lived with for thirty years. Establishing work-as-done meant walking the floor with a notebook: expensive, slow, and a sample of one week in one place. That is the constraint that has now changed.

The instrument *that changed*

Almost every step of the work now leaves a timestamp somewhere. Process mining reconstructs the process as run from those timestamps, and it does to the middle level what double-entry bookkeeping did to the organization level: it makes it countable.

The field has a settled vocabulary, fixed in its standard text (van der Aalst, 2016). *Discovery* builds the actual process model from an event log without being told what the process is supposed to be. *Conformance checking* compares what the log shows against the model the organization believes it runs: this is the one that matters for a performance analysis, and it is the one applied least. *Enhancement* extends the model with what the log also carries: durations, queues, volumes, who did what.

What that last item makes possible is the point of this section. The performer field in an event log lets you derive who handed work to whom, which turns the handoff, the step with no owner and no report, into a measurable object. The white space acquired an instrument while nobody in this profession was watching.

> **WHAT WE SEE WHEN WE RUN IT, SAID AS PRACTICE**
>
> Two findings, almost every time. The designed process turns out to describe a minority of the cases, often a small one, and nobody in the organization knew that. And the waiting time dwarfs the working time by a factor nobody would have guessed: the file spends its life in queues, and the step everyone argues about is measured in minutes.
>
> Both are ordinary findings in our work and neither is a research claim. *[practitioner observation, not a measured frequency]*

What the log *still cannot tell you*

The instrument is genuinely new and it is not an oracle. Four blind spots, and an analyst who does not know them will over-claim within two weeks.

- **It shows the route, not the reason.** A log can prove that files wait eleven days at one step. It cannot say whether that is capacity, an authorization rule, an incentive that rewards the wrong queue, or a person who is doing four jobs. Number 09 is that question, and it is still answered by talking to people.
- **Variation is not automatically waste.** Hundreds of distinct pathways through a hospital may be hundreds of appropriate clinical decisions. Counting variants is description; calling them waste is a judgement that needs its own evidence.
- **You cannot detect a deviation from a design that was never specified.** Where the intended process exists only as folklore, conformance checking has nothing to check against, and the first finding of the analysis is that absence.
- **The log records what the system was told, not what happened.** A step closed in bulk on Friday afternoon reads as forty completions in one minute. Every event log carries this kind of artifact and it has to be examined before anything is concluded from a duration.

AND ONE NOTE ON THE STATE OF THE EVIDENCE

There is a substantial published literature on process mining and very little of it is synthesized, so this paper describes what the instrument does and quotes no figures from individual deployments. What can be said without a study is that the instrument exists, that it is now cheap, and that the constraint it removes, the cost of seeing the middle level, was the reason that level went unexamined for thirty years.

Four tests
that locate a level

The level is a finding, not a preference, and it can be established cheaply. Four tests, in the order we run them. None takes longer than a day, the first one is free, and each either settles the question or names the next test to run.

THE TEST	WHAT IT COSTS	WHAT IT SETTLES
1 · The distribution, by performer	Nothing. The figures sit in the report that raised the alarm.	A narrow band (nobody fast, nobody slow) points away from the job level: whatever holds the result down reaches every performer at once, so it sits in what they share. A wide spread under the same conditions keeps the job level in play. Number 06 reads distributions in full; here the distribution serves one narrower purpose.
2 · One file, end to end	A day, with a notebook and the case system.	Working time against waiting time. A file that spends its life in queues has a process-level problem whoever handles it, because no remedy aimed at a person shortens a queue that person cannot see.
3 · The ceiling	An hour of arithmetic.	Assume every performer worked at the exemplary level and recompute the result. If the gap still stands, no intervention aimed at people can close it, and the search moves up a level. Number 07 supplies the exemplary figure this test needs.
4 · The event log	A day, where timestamps exist.	The route as run, the variants, and where the time actually sits. Section 05 describes the instrument and its limits; what it shows is a description of the problem, not yet a cause: the why still comes from number 09.

Four tests in rising order of effort. Each is cheaper than the one after it, so a cheap test that settles the level makes the expensive ones unnecessary.

The order is the discipline. In our practice, most proposals for training arrive before the first test has been run, on figures that were already in the room. The case that follows runs all four tests on one measured gap, and the proposal on the table did not survive the first two.

The insurer
whose handlers were not slow

A specialty insurer, personal injury claims. Settlement time averaged 26.4 working days over four quarters against a desired 15. Everyone agreed the department was under-performing, and the proposal on the table was a case-management training program for the eighteen handlers.

THE STEP	WHAT IT FOUND
The distribution, by handler	Uniform. Every one of the eighteen sat between 24 and 29 days. Nobody was an outlier in either direction, which is the pattern that says look at what they all share.
The single file	One claim, followed end to end. Working time, someone actually doing something to the file, 3.1 days. Waiting time 23.3 days. Total 26.4.
The event log	4,812 claims and 31,900 recorded events over the same four quarters. The quality system described a seven-step process. The log contained 412 distinct variants, and 38 percent of claims followed the designed route.
Where the waiting sat	Of the 23.3 waiting days, 14.6 sat in a single queue: files awaiting the opinion of the medical adviser.
The adviser	One person, contracted for three days every two weeks, with no view of the queue length and no rule about which files needed an opinion at all. Every file was sent.

Five steps, four of them cheap. The training proposal was still on the table when the second step finished.

Note what the analysis did not establish and did not need to. Nobody measured whether the handlers were good at case management. The question was never reached, because the ceiling test settled it: at 3.1 days of working time against a 15-day requirement, even a handler working twice as fast would have moved the total by a day and a half.

What was changed, *and what was left standing*

Two changes, neither of them a program, and both of them at the level the cause actually sat at.

The first was a triage rule. The claims team and the adviser agreed, in one afternoon, a written test for which files genuinely required a medical opinion: contested causation, an incapacity claim beyond twelve weeks, or a prior condition on file. That took the adviser's volume from every claim to 41 percent of claims. The second was capacity: the contract went from three days every two weeks to five. Neither change touched a handler.

MEASURE	BEFORE	AFTER TWO QUARTERS
Wait for the medical opinion	14.6 days	5.4 days
Total settlement time	26.4 days	17.2 days
Against the desired level of 15	11.4 days short	2.2 days short

26.4 − 14.6 + 5.4 = 17.2. The arithmetic is deliberately visible, because a claimed improvement that cannot be reconstructed is not a finding.

And the honest ending. The remaining 2.2 days sat in a second queue: a fraud screen that the same analysis had identified and costed, and that nobody funded that year. The department reported the gap as closed, which it was not, and the analysis had said in advance exactly where the residue would sit. That is the most a diagnosis can do: it cannot make the decision, which is the limit number 02 states about the discipline as a whole.

What this *does not claim*

This paper argues that the middle level is under-examined. That is a claim about attention, and it is worth separating carefully from four claims it does not make.

- **No number is being claimed for how often causes sit at each level.** There is no study that partitions established causes across organization, process and job. The argument here is about the order of examination, which holds whatever that distribution turns out to be, and section 03 says what research would settle it.
- **Neither synthesis is about a process crossing four departments.** The boundary-management meta-analysis is about teams; the integration meta-analysis is about firms and is correlational. Together they support the direction of the argument at two different altitudes, and neither of them proves it.
- **Process mining is an instrument, not a diagnosis.** It establishes where the work goes and where it waits. Everything about why remains to be established by the next question, and a report that stops at the mined picture has produced a very good description of a problem.
- **Locating the level does not make the fix easy.** A process-level cause requires two departments to change something, which is political work with its own literature and its own failure rate. Number 10 takes the choice of remedy; nothing in this paper makes the negotiation shorter.

SUMMARY

Three levels, and a remedy reaches only one of them. The middle level produces the result, has no owner and no end-to-end report, and until recently could only be examined by walking the floor. That constraint has lifted. What has not changed is the reflex: when the middle level is invisible, the analysis reports on the level it can see.

RUN IT YOURSELF

Six lines,
on the gap in front of you

Half a day, on the gap you established in question two. Line three is the one that cannot be answered from a desk, and line six is the one that changes the budget.

1 · The distribution, by person

Uniform, concentrated or two clusters? You cut this in number 06. Write the answer here before you look at anything else.

2 · The intended process

Find the document. How many steps, and when was it last revised? If it does not exist, that is finding number one and it is not a small one.

3 · One file, end to end

Follow a single unit of work. Record every handoff and every wait, with times. Working time and waiting time, separately.

4 · The biggest queue

Which single wait accounts for the largest share of the total? Who owns the step on the other side of it, and do they know how long the queue is?

5 · The transplant question

Swap your exemplary performer with a standard performer elsewhere. Do the results follow the people or stay with the places?

6 · The ceiling

If everybody performed at the exemplary level, what would the result be? Is that still short of the desired level? If yes, no intervention aimed at people can close this gap.

WHERE THIS PAPER SITS

Two problems, *one method, sixteen papers*

	BUSINESS PROBLEM	TRAINING REQUEST
Owned by	The manager or director who owns the result.	L&D and the HR business partner.
What it prevents	Spending on the wrong intervention: a system, a reorganization, a hire.	Spending on a program that cannot reach the cause.
What it returns	The fix that moves the number, chosen on evidence.	The budget and the working hours that were not spent.
Method used	The same three analyses and five questions, which is why one capability serves both, and why the diagnosis no longer has to be bought in case by case.	

One method, two returns, and a standing capability instead of a standing consultancy line.

Two models carry the method. Gilbert's six conditions ask what kind of cause; Rummler and Brache's three levels ask where it sits. Both are sound and this series uses both, and it leads with Gilbert, because his unit is the accomplishment, which is what an organization already counts, and because his six conditions are the categories an organization can actually change.

Where the method *comes from*

Two further pieces close the set and carry no number, because each points at all the numbered papers. Before the guide sits Arets L&D Business Models, the destination piece for the people who receive the requests: four models a learning function can run, and the route from HR-driven to business-driven. Last sits Solutions that cannot work, the field guide, ordered by the sentence that lands on your desk: seventeen forms a performance problem arrives in, each with an address in one of the numbered papers. Three orders in all: the destination you choose, the question you are on, and the form that landed on your desk.

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.

About the cases. Every worked case in this series comes from our own practice. Names are changed and details altered so that no organization is recognizable, and the arithmetic reconciles. That is said once, here, rather than under each case: client work is confidential, and a case that named anybody would be a breach rather than a proof.

THE SERIES

Sixteen papers, *six parts*

Number 01 is the map and number 02 is the foundation. From number 05 on, every number is depth on exactly one step of the method. The filled block is the paper you are holding; the two unnumbered pieces close the set.

I · WHAT HPT IS

The discipline, where it came from, and why the reflex it replaces survives.

- 01 The performance improvement process
- 02 Human performance technology
- 03 Environment before capability
- 04 The system and the performer

II · THE BUSINESS ANALYSIS

Question 1.

- 05 Business needs and performance needs

III · THE PERFORMANCE ANALYSIS

Questions 2 and 3.

- 06 Gap analysis
- 07 The exemplary performer

IV · THE CAUSE ANALYSIS

Questions 4 and 5.

- **08 Cause analysis: the three levels**
- 09 Cause analysis: the six conditions

V · THE DECISION

From cause to solution, and what counts as proof.

- 10 Intervention selection
- 11 Learning metrics and business metrics

VI · THE FUNCTION

The request, the money, the implementation, the assistant and the value.

- 12 The solution request
- 13 The cost of non-performance
- 14 Implementation and maintenance
- 15 Performance augmentation
- 16 The Value Flip Calculator

Arets L&D Business Models · the destination piece, after 16
No number: the four positions a function can hold, and the move the sixteen papers make simple.

Solutions that cannot work · the field guide, at the end
No number: it points at all sixteen. The same material, ordered by the sentence that lands on your desk.

Business problem: a result is moving the wrong way · Training request: a solution has already been chosen · Both doors: after the decision

REFERENCES

Where every claim *in this paper comes from*

Desmedt, M., Ulenaers, D., Grosemans, J., Hellings, J., & Bergs, J. (2021). Clinical handover and handoff in healthcare: A systematic review of systematic reviews. *International Journal for Quality in Health Care, 33*(1), mzaa170. https://doi.org/10.1093/intqhc/mzaa170

Gilbert, T. F. (1978). *Human competence: Engineering worthy performance.* McGraw-Hill.

Leicht-Deobald, U., Backmann, J., de Vries, T. A., Weiss, M., Hohmann, S., Walter, F., van der Vegt, G. S., & Hoegl, M. (2025). A contingency framework for the performance consequences of team boundary management: A meta-analysis of 30 years of research. *Journal of Management, 51*(2), 704–747. https://doi.org/10.1177/01492063231206107

Leuschner, R., Rogers, D. S., & Charvet, F. F. (2013). A meta-analysis of supply chain integration and firm performance. *Journal of Supply Chain Management, 49*(2), 34–57. https://doi.org/10.1111/jscm.12013

Rummler, G. A., & Brache, A. P. (2013). *Improving performance: How to manage the white space on the organization chart* (3rd ed.). Jossey-Bass. (Original work published 1990)

van der Aalst, W. M. P. (2016). *Process mining: Data science in action* (2nd ed.). Springer. https://doi.org/10.1007/978-3-662-49851-4

The three books named earlier in this paper are listed in full in number 02, so that they do not recur in every reference list.