Future Processing

Our path from AI adoption to AI transformation

Best practices from our method, measurements and results

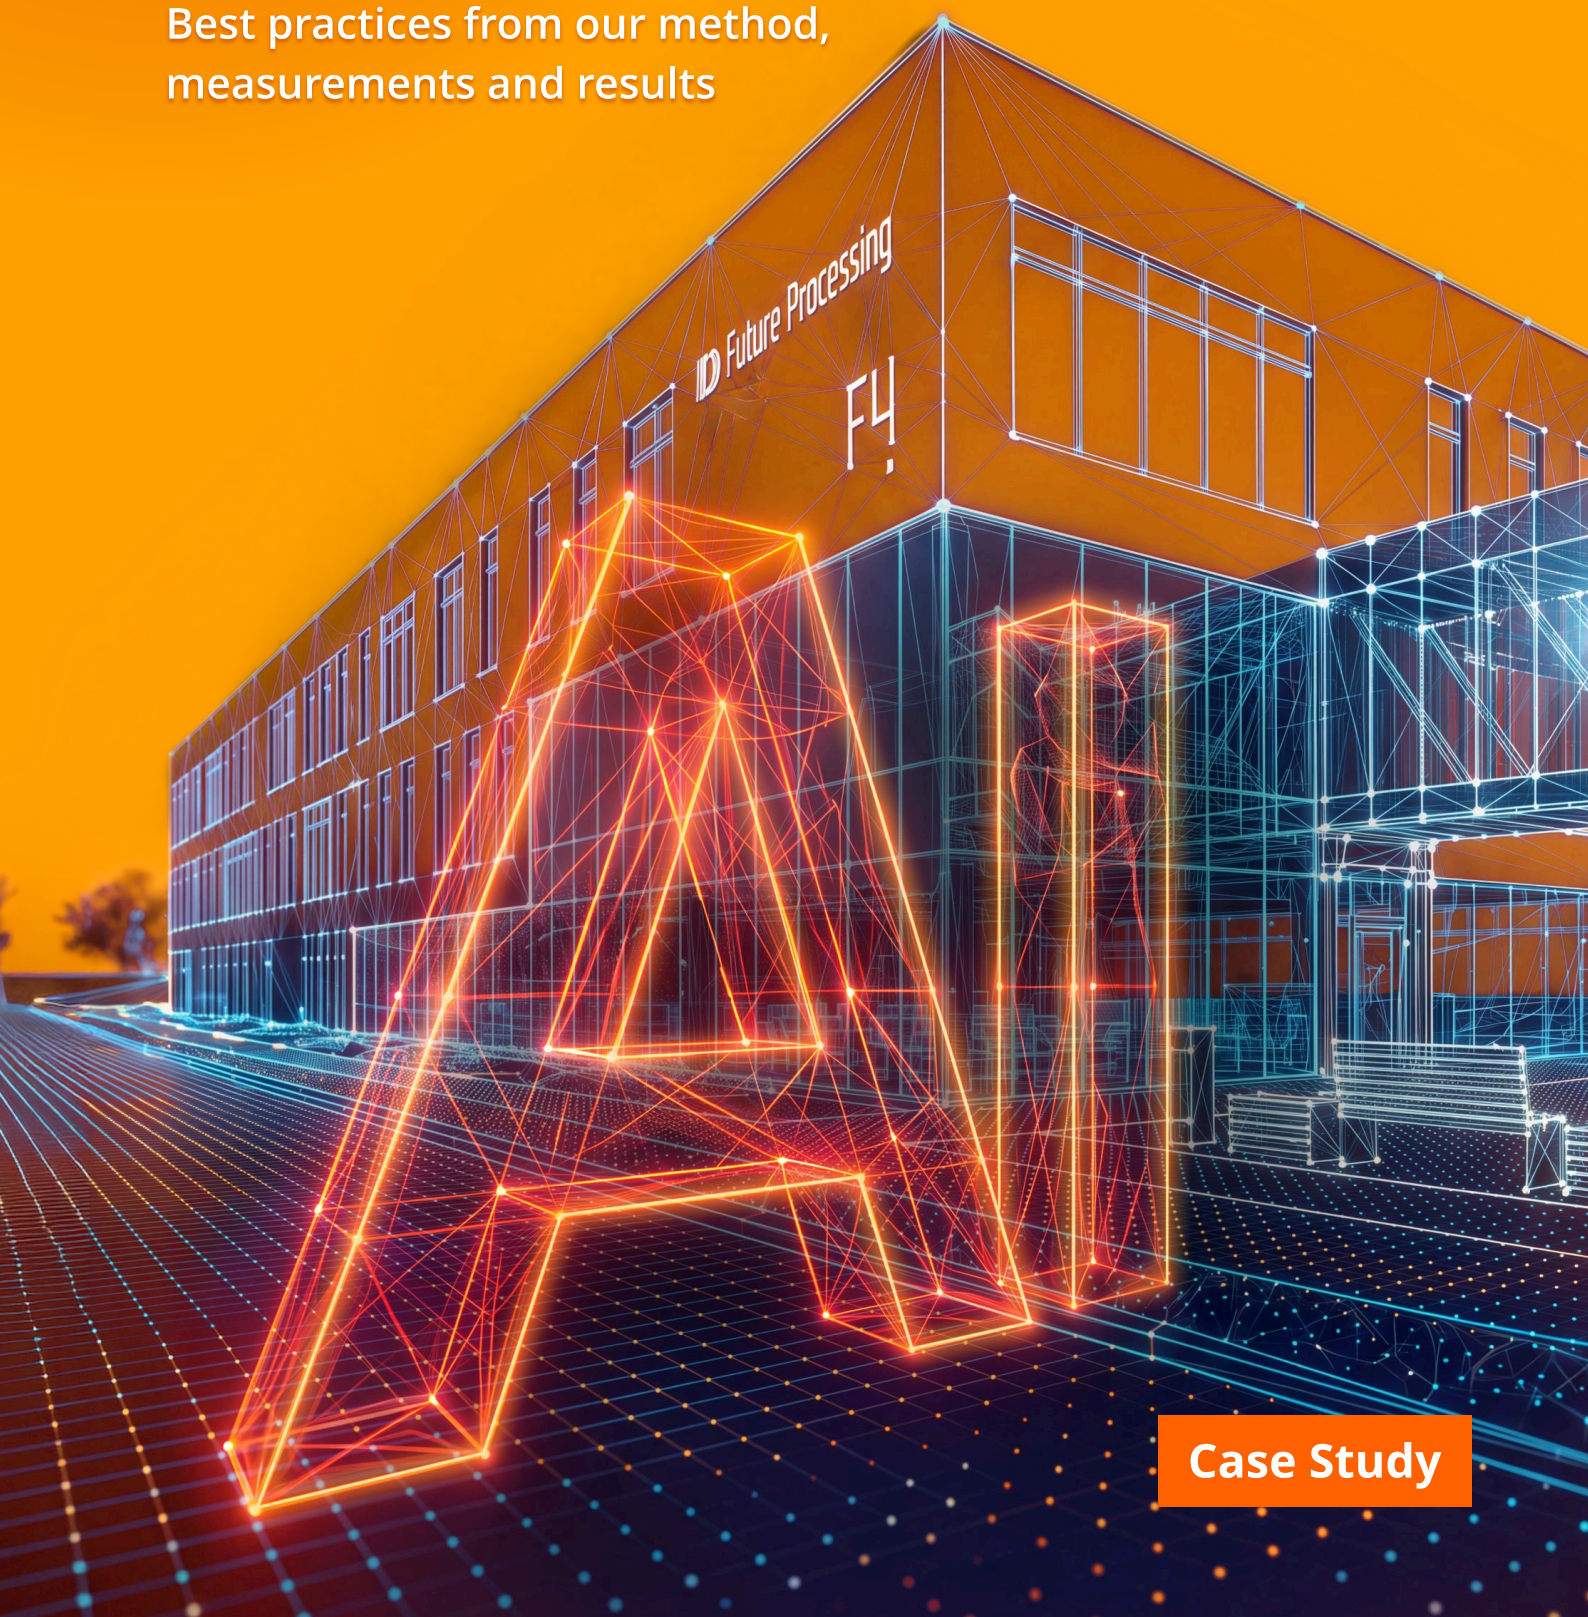

Case Study

Contents

The case in brief **Everyone uses AI. Using it as a team is the harder part** 3

Part one **Where did we stand with AI, and how did we know?** 7

Part two **From adoption to transformation** 14

- Stage one Point adoption 16
- Stage two Wide adoption 20
- Stage three Shared AI competence 23
- Stage four Organisational resilience 33

AI-native delivery **How we rebuilt software delivery around AI** 40

Part three **Where to start, and what to do next** 47

About us **About Future Processing** 51

How we measured **A note on method** 52

The case in brief

Everyone uses AI. Using it as a team is the harder part

In our first company-wide survey, 81% of our people said they used AI tools. In the same survey, the same people gave 34 out of 100 to how they used AI as a team. The first number describes habits. The second describes an organisation, and the organisation had not yet changed.

81%

of our people said they used AI tools

34/100

the score the same people gave to how they used AI as a team

Employee survey, first wave, n=446, a 66% response rate.

The distance between those two numbers is where most companies stand with AI today. This case study describes how we closed a large part of it in our own company over fifteen months, and what we would do the same way again.

The challenge

In early 2025, every company that wanted to keep up with the technology faced the same question. The business opportunities in AI were becoming clear. What nobody in the industry knew yet was how to turn them into results that would show in a company's financial statements.

Some companies went all in, investing early and widely in the hope that a few of the bets would pay off. Others started small, with dedicated teams running internal pilots and trying to find a business case among countless possibilities.

We chose a path of our own. It took something from both, and it rested on different principles. We treated ourselves as patient zero: we laid the groundwork first, built the organisation's capabilities and set up the measurement before anything else, so every result in this case study is one we measured.

The board of Future Processing concluded that AI would change how software is built, and decided that the company as a whole had to transform the way it operates, with a plan, rather than leave AI to individuals. Individual use was already spreading on its own.

A few engineers, analysts and consultants were far ahead, and what they learned stayed with them: it was hard to repeat, share across a team or use on a client's project, and it left when they did.

> *AI is changing what every one of us does at work. At Future Processing we decided to meet that change as a whole company and to treat it as a chance to grow: to learn the new way of working ourselves, together, so that our people and our clients can rely on it.*

— **Michał Sztanga**
CEO

For a software company that delivers for clients in regulated industries, the business challenge was to make AI part of how every team works, in the same way and to the same standard, without lowering our quality and security standards.

The approach

The goal we set was a systemic upgrade of our people's AI skills, and a redesign of our internal processes around the way people work with AI.

By systemic we mean that everyone works with AI in a standardised way, in their projects and in their teamwork, and stays open to experiment. We set no cost or time-saving targets at the start. We set skill targets, and we measured them.

Before we started, we set ourselves boundaries. The adoption frameworks on the market did not fit what we wanted to see and measure, so instead of adapting one, we built our own. Mike Zamarski, our AI Strategy Advisor, designed the four-point AI Organisational Resilience Framework: four stages, each defined by the main challenge a company faces at that point of adopting AI.

The framework deliberately has no separate technology layer. The technology was changing faster than any plan could follow, so we wanted a company that can take up a new tool as soon as it proves useful, and retire one when something better arrives.

To react that quickly, we first needed a way to see where we stood. So we designed a company-wide study of around ninety AI-related factors, taken once as a baseline before any systemic change and then repeated in cycles, so that the difference between two readings shows what AI is changing in the company.

Most of the numbers in this case study come from it.

STAGE	WHAT WE DID
1 · Point adoption	A board decision, a baseline survey of the whole company, and a plan that fixed the goals and the budget and left the route open
2 · Wide adoption	One funded route to the tools with the rules attached, a one-page policy with a review date, and an open invitation to experiment
3 · Shared AI competence	A skills programme for every role, mentors and a champion in every project, experiments written into team goals, and a quarterly maturity assessment
4 · Organisational resilience	New standards of work in our key business processes, from lead qualification and presales to delivery, a library of proven methods, replaceable tools, and a business model built more on results than on hours

The study is an employee survey, a standard instrument of market research, and 66% of our people answered its first wave. Survey answers are self-reported, so we read them alongside a quarterly maturity assessment scored against a rubric. The note on method at the end of this case study describes both.

The results

The skills moved where we had aimed them: at team level. In the first survey, our people gave 34 out of 100 to how they used AI as a team. By the second, in April 2026, the score was 52.7, and no other measure in the study moved as fast. It moved together with how people see their projects: the more a team works with AI together, the more its members say AI helped the project succeed. Individual use on its own shows no such link.

Both survey readings came before our shared standards of work were in place, so the effect of those standards has yet to show. The quarterly maturity assessment, in which the head of each competence centre scores the centre against a fixed rubric, already shows the skills building up: people, practices and governance rose in every one of the three readings so far.

MEASURE	MOVEMENT
AI used as a team, employee survey, out of 100	34 to 52.7
The gap between what people wanted from the company and what it gave them	26 to 15
Daily use of AI, whole company	doubled
People capability in the competence centres, maturity assessment, out of 5	2.56 to 2.89 over three quarterly readings
AI-supported practices defined for each role, maturity assessment, out of 5	2.11 to 2.56
Delivery methods approved for reuse, each proven in more than one client project	6 to 18

Employee survey: wave one n=446, 66% response; wave two n=480; self-reported. Maturity assessment: scored quarterly by the head of each competence centre against a five-level rubric (Stage three). Delivery methods: the use-case library

Our key business processes now run on new standards of work. In delivery, AI-native delivery became the standard way our teams work, and in sales, presales and contracting the work was redesigned in the same way.

Client projects show what the new way of working does to a project. For an international strategy consulting firm, we delivered the agreed scope with more than 30% less effort than our historical benchmark. For Fortaco, we had a fully working prototype, ready for the first tests of its SAP integration, within two weeks, and it changed the way we work with the client.

Business results at the scale of the whole organisation are coming next, and the order is deliberate. Skills come first, then the capacity and a standardised way of working, and then efficiency and return, as the new way of working spreads across client projects and more contracts pay for outcomes. That spread is under way.

Client code and data stayed under the same security rules as before AI arrived, and our quality rules did not change.

Our information security management system has been certified to ISO 27001 since 2017, and every client's contract comes first: where a client restricts AI, the client's rule prevails on their project. What is new is the mechanism that enforces the rules, and each stage below describes it.

Part one

Where did we stand with AI, and how did we know?

AI adoption, maturity and transformation: how we tell them apart

When we started measuring how AI was changing Future Processing, it became clear that the change has several sides, and each set of numbers makes sense only in its own context. Remigiusz Tunowski, our BI Consultant, proposed the distinction we use throughout this case study.

Adoption is uptake: whether AI is used, by whom and how often. Transformation is redesign: whether the company does its work differently because of AI. “Twenty per cent of European firms use AI” and “AI has transformed twenty per cent of European firms” are different sentences, and most headline figures in circulation are the first being read as the second.

Between the two sits maturity: how good the skills and practices behind the use are. We keep the three words apart and read each with a different instrument.

TERM	WHAT WE MEAN BY IT	HOW WE READ IT
Adoption	Uptake: who uses AI, and how often	The employee survey, two waves
Maturity	Capability: how good the skills and practices behind the use are, scored against a rubric	The quarterly AI Maturity Assessment, per competence centre (stage three)
Transformation	Redesign: the company does its work differently because of AI	Not scored by the survey. It shows in the methods proven on client work, in the client projects and in the processes and roles that changed

Why we started with a survey

The survey was the first thing the transformation decision produced. Once the board had decided that the whole company would go through an AI transformation, the CEO commissioned a cycle of strategy workshops to turn that decision into a plan. The first initiative the workshops agreed on was a baseline: **before building anything, measure where the whole company stood.**

A baseline served the plan in two ways. It showed us where to start, and it gave us something to measure progress against, because the survey was designed from day one to be repeated once the programme was under way.

Market data could serve neither purpose. Benchmarks show whether firms use AI, not how far it has been adopted inside any one of them. How widely our teams used it, what our people thought of it and what they expected from us were questions only our own people could answer.

> *Once the decision to transform had been taken, we knew we had to ground it in business reality, in hard numbers. A baseline works best at the very beginning, so we captured ours before any systemic change took place. It took real effort at first, and it pays off every time we launch a new initiative, because we can track its effect on our organisation.*

— Mike Zamarski
AI Strategy Advisor
the mind behind Future Processing's AI strategy and the author of the four-point AI Organisational Resilience Framework

WORTH COPYING: AI GOVERNANCE

Measure before you start

The move

A company-wide survey taken before the programme had delivered anything, designed from the first day to be run twice.

How you would use it

Run yours now, before the next decision, with a team-level question in it, and fix the date of the second wave at the same time. Two readings tell you what your own actions did, which no market figure can.

Watch for

Reading a single number without context. Compare like with like, and read each result alongside what changed between the two readings.

So we asked all our employees how they used AI, what they thought of it and what they wanted from the company. The first wave set the baseline. The second followed in the middle of the work.

How we measured. Around ninety questions went to every employee in the first wave, and 446 people answered, a 66% response rate. The second wave used a third of the questions and kept the comparable core in: 480 respondents, of whom 293 had answered the first time. A third wave is planned. Every figure in this part is what people said about themselves.

Measuring the team, not the person

We measure AI use at the level of the team, and manage it there, because **a time saving counts only when the whole chain of work moves.** When one person finishes an hour earlier, neither the rest of the team nor the client gains anything unless the next step, the review, the approval or the hand-off, starts an hour earlier too. Usually it does not, because each of those steps runs to its own rhythm.

> *When AI saves one person time, the next person in the process still expects the data two hours later. That is why adding AI to one step is not enough: the whole process has to be redesigned around it.*

— **Dominik Samociuk**
Head of Security & Compliance

The same thing happens at the scale of a project. When AI makes one part of the work much faster, the team's efficiency does not necessarily rise at the same pace, because the parts around it start to set the pace. **Turning many small gains made by individuals into a gain for the whole team is a balancing exercise, and the hardest part of an AI transformation.**

> *When one part of the process speeds up, a new wall appears somewhere else. Generate ten times more code, and if people still read every line, they stop doing anything but read. One person being faster gives us nothing. The whole team has to be faster, in every phase.*

— **Konrad Bałys**
Head of Engineering

So we measured the team. Between the two waves, team-level practice moved from 34 to 52.7, faster than any other measure in the study.

It is also linked to how projects go. Teams that practise AI together rate AI's contribution to their project's success higher (a correlation of r = 0.42), and the relationship holds after allowing for how much each person uses AI on their own. **Individual use, on its own, does not predict it.**

What we changed between the two surveys

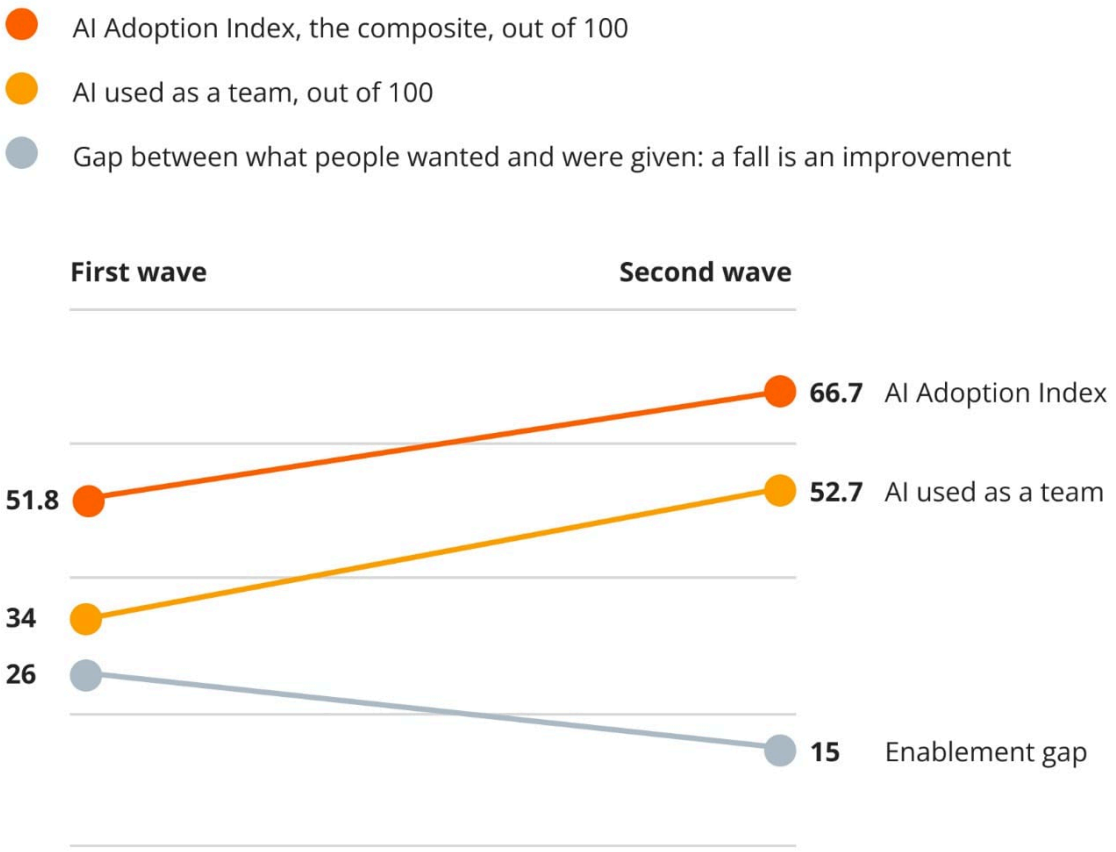

Employee survey, first wave (n=446) and second wave (n=480). Self-reported. The gap is in the same survey points and drawn so that its fall reads as improvement. Team practice correlates with how much respondents say AI helped their project succeed, r = 0.42 (Spearman, n=480), holding after controlling for use frequency.

How the numbers moved

Daily use of AI doubled across the company between the two readings. It rose in every delivery role by a factor of between 1.7 and 2.5, and the roles that started ahead stayed ahead. A programme that gives everyone the same tools lifts every role, and it leaves the gaps between them where they were. That is why the next stage wrote AI expectations for each role.

Daily use of AI by delivery role, second wave as a multiple of the first

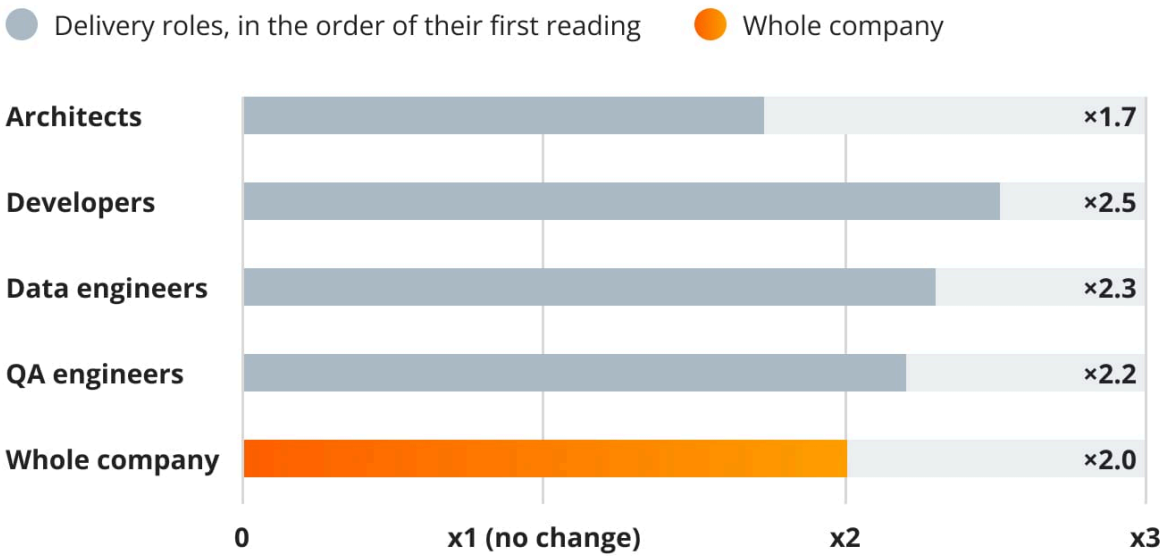

Employee survey, first wave (n=446) and second wave (n=480). Self-reported. Printed as growth on the first reading; the roles are in the order of their first reading, which did not change

WORTH COPYING: PEOPLE

Read your survey as a ladder

The move

We sorted the answers into four steps, each harder than the one before: people know about AI; they try it now and then; they use it regularly on their own work; their team uses it together, in a way it repeats and shares. Our 81% sits on the middle steps. Our 34 out of 100 describes the top one.

How you would use it

Sort your own results the same way. A high usage figure shows that people have reached the middle steps. It says nothing about the top step, and the top step is the one that changes how a company works.

Watch for

A usage percentage reported to the board as if it described the top step.

The gap between what people wanted from the company and what it gave them narrowed from 26 to 15. People wanted more than before (their desire to learn rose from 69 to 75 out of 100), and the company caught up faster: their rating of its readiness to support them rose from 43 to 60, and their knowledge of the tools from 37 to 52. What remains is mostly knowledge of the tools.

The finding worth noting is where it stayed widest: with leaders, consultants and top management rather than with the people furthest behind. These were the people expected to set the standards and open the way for everyone else, and they felt the least supported. That set the work of the next stage.

The survey found no resistance to speak of. The section designed to measure it came back close to empty in both waves, which surprised the people who built the survey. The share of people placing themselves as sceptics fell by more than half, and the share calling themselves enthusiasts or pioneers rose by half. People were willing to grow; the work of the year lay in giving them the right support to develop the skills.

Where AI went deepest

Some of the deepest adoption in the company is in roles that are not technical at all. In the second survey wave, two of them scored among the highest of any role and well above the company average: marketing and communications, and consulting. The score covers all four steps of the survey ladder, so it shows how far AI has gone into the work itself, not only how often people open the tools.

Both roles spend much of their time on preparation: the research, analysis and groundwork behind a campaign, a recommendation or a discovery workshop. That is where AI takes the most off people's hands, and it leaves more time for the part clients pay for. It also helps explain why consultants were among those asking most of the company: the further people go with AI, the more they want.

> *In marketing, communication and consulting, a lot of time used to go into groundwork: collecting data, compiling reports, sorting through research, writing up workshop notes. AI now takes much of that over, so people can spend their time on strategy and ideas, where they add real value. The shift has a clear cause. Clients no longer want to pay for keeping processes running. They want to pay for groundbreaking ideas that deliver business outcomes.*

— **Filip Sitko**
Head of Marketing

What a survey cannot tell you

A survey records how people assess themselves at one point in time, so its results are subjective. Whether the way a company works has changed shows in the work, across the whole chain of it. For that we used two more instruments: the quarterly maturity assessment, which scores skills and practices against a rubric (stage three), and the use-case library, which holds the methods proven in client work (stage four).

KEY TAKEAWAYS FROM PART ONE

- **Measure the whole company before you build anything**, and design the second reading at the same time. A single reading tells you where you are; two tell you what you did.
- **Put a team-level question in the survey**. Usage alone will flatter you.
- **Do not read survey results as pass or fail.** Read each one in context, alongside what changed between the readings.
- **Keep adoption, maturity and transformation apart**, and do not let anybody read a usage percentage as a transformation score, including your own board.

Part two

From adoption to transformation

We describe AI adoption inside a company as a sequence of stages, each named after what the company has reached by the end of it. In **point adoption**, individuals use AI on their own. In **wide adoption**, the whole company has access to the tools and works under shared rules. In **shared AI competence**, AI skills become part of each role, starting with delivery, and teams apply them to a common standard. In **organisational resilience**, the company can absorb whatever AI brings next.

The first two stages are adoption: AI moves from individual experiments into parts of the company's processes. The last two are transformation. People's skills change first, and with them the way teams use AI in their daily work; then the way the company itself operates.

The stages: what each means, and what we did

WHAT IT MEANS		WHAT WE DID
Individuals use AI on their own	**Point adoption**	A decision at the top, and a baseline reading
Everyone has access, with rules made for AI	**Wide adoption**	One route to the tools, rules attached, and an invitation to experiment
Every role has the skills, and teams use them alike	**Shared AI competence**	A skills programme for every role, and a quarterly maturity score
The company absorbs each new change in AI	**Organisational resilience**	New standards of work in key processes, replaceable tools, results-based model

Our own description of how a company moves from AI adoption to AI transformation.

STAGE	WHAT IT IS	MINDSET	MAIN RISK
1 · Point adoption	People use AI tools on their own initiative, inside the company's existing security rules, with nothing shared, repeated or measured.	tools as experiments	value stays with individuals; effort is duplicated
2 · Wide adoption	Use spreads across teams and processes. A sanctioned route to the tools, rules designed specifically for AI use, and an open invitation to experiment are built to match it.	tools as processes	inconsistency between teams; expectations run ahead of results
3 · Shared AI competence	AI skills are built systematically in every role, so that teams work with AI in the same, standardised way and keep experimenting.	skills and shared standards	skills learned but not applied in projects; a standard that only ever tightens, with nothing that keeps it open to new methods
4 · Organisational resilience	The company absorbs and adapts to AI at scale, in its delivery model, its tools, its governance and its commercial model, without losing control, compliance or culture.	culture and systems as the foundation	treating the new operating model as finished while AI keeps changing; change fatigue, when people are asked to change again before the last change has settled

How long does an AI transformation take?

Our transformation is not designed to have an end point. It took fifteen months to get from the baseline reading to the first milestone we could measure: a delivery standard in use on client work. The stages took very different lengths of time, and each one carries on as a programme of its own.

Point adoption produced a decision and a plan over a few months, and it continues as our AI Experiments Programme. Wide adoption was in place within a single quarter and is now carried by the AI Masters initiative. Shared AI competence took half a year of goals, roles, experiments and quarterly measurement to get going, and it continues as the AI-Boosted Delivery programme.

Organisational resilience is continuous work with no end date: tools and regulations keep changing, and resilience is what lets us adapt when they do.

Each stage asks for a different kind of solution, and what worked in one does not carry into the next. A programme stalls when any of its stages fails to deliver, and keeping them in balance takes conscious effort and continuous investment.

The order matters too. A company cannot start at Shared AI competence: mentors, champions and a skills programme spread a way of working, but the decision about what that way of working is for belongs to point adoption. Without that decision, a skills programme has nothing to standardise towards.

Stage one

Point adoption

Everyone's already using AI

Doesn't that mean we've adopted it?

It means the tools are in use, and it says nothing yet about the company. Ours began with a handful of engineers, analysts and consultants testing the models on their own, under the security rules and client confidentiality terms that already applied to all of our work. Our engineers also formed a hypothesis group of around thirty people to test what AI could be used for. From April 2024 anyone in the company could also join AI Coffee (stage three).

One of the people behind our AI strategy, Mike Zamarski, was researching how AI affects the way people work, and he took his conclusion to the top of the company.

> *Long before our AI transformation began, I came to a conclusion that hardly anyone was voicing at the time: the hard part of the journey would be human adoption and organisational inertia, more than the technology, and our transformation had to allow for that.*

— Mike Zamarski
AI Strategy Advisor
the mind behind Future Processing's AI strategy and the author of the four-point AI Organisational Resilience Framework

If AI is an organisational challenge, no amount of individual use moves a company forward, and the question becomes one that only the top of the company can answer.

The decision to run a transformation

At a management review, the board decided that the company would commit fully to AI. Future Processing would run an AI transformation as a programme of the whole company, with a plan, a budget, a sponsor for each stream of work and a checkpoint at the end of the year.

The programme's goal was set with it: **a systemic upgrade of our people's AI skills, so that teams could work with AI in a standardised way, and a redesign of our internal processes and workflows, so that they reflect how people work once AI is part of the job.**

> *We decided that AI would be a change for the whole company, with a plan, a budget and a goal everyone could understand: our people would learn to work with AI together, to one standard, and the way we work would change to match.*

— **Michał Sztanga**
CEO

THE SECURITY QUESTION AT THIS STAGE

How is client data protected while AI use is still individual?

By the rules that already govern all of the work. Our information security management system has been certified to ISO 27001 since 2017, and every client contract carries its own confidentiality terms. Both applied to every tool, AI included, from the first day.

Our security function's first step was a list of approved AI tools, assessed one by one. As use grew, the team saw that the same tool carries one level of risk for an internal meeting and another for a client's product roadmap, so the rules grew from approving tools to also governing how AI is used, on what and for whom. That change runs through the stages that follow.

For a reader at this stage. Start from the rules you already have, because they already cover much of AI use, and plan to add the governance of how AI is used to the approval of tools.

Who was in the room

We put the people who would have to do the work in the room, the cautious ones included. The CEO commissioned a cycle of internal strategy workshops, with around thirty people chosen to represent the main streams and departments of the company and to bring different views to the table.

One assumption was set at the start: a company does not move at one speed. Each part of the company was worked through in a conservative, a balanced and a revolutionary version, and the plan that came out has different speeds inside it. That is what let the slower parts of the company sign it.

WORTH COPYING: AI GOVERNANCE

One direction at different speeds

The move. We worked through the plan in three streams. Sales could take more risk and move faster. Delivery put quality first. People and Culture protected the company's heritage and values. For each stream we wrote a conservative, a balanced and a revolutionary scenario, and the final plan combined them, with a different pace for each stream.

How you would use it. Let the parts of your company that must protect quality or people plan more slowly than the parts that can take risk, and write both speeds into the same plan, so that the slower parts can sign it.

Watch for. A plan that asks everyone to change at the same speed. The parts of the company that cannot keep up might read it as a big push.

What went to the board

The workshop cycle produced the decisions that shaped the year: measure the whole company before building anything, and design the second reading at the same time; make the change in our own company first; give delivery a programme and a mandate of its own.

The board then received a plan with three requests:

- approve the governance model,
- choose a direction for the infrastructure,
- approve the roadmap.

Progress would be checked at the end of the year.

The plan set the objectives and the budget. It did not lay down a route to them, by design: with the tools changing from month to month, any route fixed for a year would have been wrong by the second quarter. We reviewed direction against the objectives every quarter, and spending moved within the budget to follow it.

WORTH COPYING: AI GOVERNANCE

Report the programme by what it delivers

The move. Every task in the programme carried two labels. One said who was responsible for it: the stream and its named owner. The other said what it was for: the initiative it served, named after the result it was meant to bring.

How you would use it. Manage the work through the owners, and report it to the board through the results. The board then discusses what the programme achieves.

Watch for. A programme plan that lists owners and tasks but no purposes. It shows the board how the work is organised, and nothing about why it is being done.

When we started the AI transformation, we fixed the goals and the budget, and left the route to them open. The tools were changing from month to month, so we corrected the route as we learned, every quarter.

— **Daria Polończyk**
Head of Delivery

KEY TAKEAWAYS FROM STAGE ONE

- **Decide at the top, and say what the transformation is for.** Ours was a systemic upgrade of our people's AI skills, and processes redesigned around the way people work with AI.
- **Put the people who will do the work in the room where the plan is made**, the cautious ones included, and let different parts of the company plan at different speeds.
- **Fix the objectives and the budget, and review the route every quarter.**
- **Measure before you plan the next steps.** The starting point exists only once.

Stage two

Wide adoption

If we give everyone AI tools, how do we keep control of what goes into them?

Stage two is where AI stopped being a private habit and became something the whole company could use, safely and in the open. The pieces had to arrive together: one funded route to sanctioned tools, a first AI policy and the start of AI governance, and an open invitation to every team to experiment. Tools without rules would have created risk, and rules without tools would have pushed people towards tools nobody could see.

How people got the tools

Every request for an AI tool went through one channel, built by a cross-functional taskforce. Approval took at most two working days, and budget and usage were tracked from the day a licence was granted.

Most of our people work on client projects. A tool paid for from one project's budget would have been available only to the people on that project, and only for as long as the project lasted. So the company paid for the licences centrally, and everyone had the same tools to learn with, whichever project they worked on.

> **DEFINITION: SHADOW AI**
>
> AI tools used for work that the company cannot see: personal accounts, free versions, browser extensions. The risk lies in what nobody can see, namely what goes into these tools, what comes out, and under which terms. It usually grows where the approved route is slower than the unapproved one.

The reason for the speed was visibility. When the approved way is slow, people find their own, and shadow AI grows: tools in use for work that the company cannot see or govern. A fast approved route is the most reliable way to prevent it.

Interest and use turned out to be different numbers. Interest in the tools was high from the start, but a month after the licences opened, actual use was well below it.

Many people who wanted a tool had not yet made it part of their work. That gap is normal, and it is useful to see early, because it shows where people already have access and need help to use it. By the second survey wave, seven tools were in active use by more than 30% of people, where at the first reading only one had reached that level.

What is AI governance, and do we need it before we scale?

AI governance is the set of decisions and rules that determines who may use which AI tools, with which data, who owns the results and who is accountable for the risk.

A company needs a first version before AI use scales. That first version can be short, as long as it is clear, names its owners and carries a date for its review.

We published our first AI policy at the same time as the access route opened. It covered access, data handling and intellectual property, it took two minutes to read, and it carried the date on which it would be revised.

It was built on the certified information security system we already had, and on delivery that already had to meet regulatory requirements, among them GDPR and, where they apply to a client or a project, the AI Act, NIS2 and DORA. Spending on AI was monitored from the start and reported to the company's financial control team.

The first version named the sanctioned tools and said what may go into each. The planned review changed the question. The same tool carries one level of risk for an internal meeting and another for a client's product roadmap, so the second version of the policy assesses both the tool and how it is used.

To make that workable day to day, we built a Compliance Assistant: a conversational tool that gives anyone a quick answer to a compliance doubt, once they describe what they want to do with AI, on what data and for whom (stage four shows how it works).

THE SECURITY QUESTION AT THIS STAGE

What may go into a tool, and who owns what comes out?

The policy answered both before the licences opened wide. Access and policy arrived together, so nobody had a tool before the rules for it existed. The policy named the sanctioned tools, what may go into each, and who owns the result. We monitored usage and spend from the first day, and each rule had a named owner.

On a client's project the contract remained the first check: where it restricts AI, that restriction wins, whatever our own rules would allow. The principle behind all of it is the one we now put in front of clients: AI risks known, managed and documented.

“AI is like a hammer: a fine tool, as long as you know how and when to use it.”

Dominik Samociuk, Head of Security & Compliance

For a reader at this stage. Publish the first policy together with access, keep it to a page, and date its review.

Who owns what AI produces

Who owns what a supplier builds in a client's project is set by the contract, as it was before AI. The use of the client's code and data with AI stays within the scope and safeguards agreed in that contract. Training a model on them and using them as context for a task (grounding) are assessed separately, and approval for one does not imply approval for the other.

The same question arises inside a company the day the tools arrive. Our first policy answered it on one page: what people produce with AI in their work belongs to the company, within each tool's terms of use, and what may not go into a tool is written next to that rule.

An open invitation to experiment

With the tools came an invitation, and it went to every team in the company: delivery, and equally marketing, sales, HR, recruitment, finance and the other teams that support them. From then on everyone was encouraged to experiment on their own work, on their team's processes and in their projects, alone or together. The message was simple: try it, use it, check it.

Nothing was assigned and nothing had to be proven. Experiments were not yet part of anybody's goals, and a failed attempt taught as much as a successful one, because the point was to find out what AI could do on real work.

To keep what people learned, we opened a register of experiments that anyone in the company could read and add to. From it we picked the practices with the potential to become shared practices.

The habit of experimenting that started here is what stage three turned into a system.

KEY TAKEAWAYS FROM STAGE TWO

- **Publish a short policy together with access, and date its review.** Plan for the second version to assess how AI is used as well as which tool.
- **Fund the tools centrally, and put a service level on the approved route.** The fast route is the one people take.
- **Wire usage and spend into the approval step from day one**, and expect use to trail interest at first.
- **Put the ownership of AI-assisted outputs in writing**: in your contracts with suppliers, and in your own policy.
- **Invite every team to experiment on its own work, not only the technical ones, before you ask for anything in return**, and let a failed experiment count.

Stage three

Shared AI competence

How do we build AI capability when we can't hire it?

We grew it inside the roles people already had. AI skills were not available on the market at a sensible price, and a training campaign on its own would not have produced them.

By the start of this stage everyone who wanted a tool had one, and the rules held. What the company still lacked was skills spread evenly enough for teams to work with AI in the same way. The two survey readings showed where the work lay:

- **The most advanced people felt the least supported.** The gap between what people wanted from the company and what it gave them stayed widest among leaders, consultants and top management.
- **Every role moved, and the gaps between roles stayed.** Daily use rose in every delivery role, by between 1.7 and 2.5 times, and the roles that started ahead stayed ahead (see the chart in Part one).

> *We wanted AI maturity first. If our people know how to work with AI, and work with it in the same way, the return will follow and will speak for itself. Counting the hours saved on single tasks would have told us very little about that.*

— **Konrad Bałys**
Head of Engineering

AI as a skill in every role

The People and Culture stream had been planned at the most careful of the three speeds, and this is the stage where that choice showed. Skills were built through the roles people already had, without a reorganisation. That was a conscious choice, made with the three speeds described in stage one: we judged that changes to people and talent management would succeed only at a more conservative pace.

The principle went to the whole company before anybody was appointed to anything: AI is a skill inside every role, the way Git or English is for a developer, so there is no dedicated AI role.

THE SECURITY QUESTION AT THIS STAGE

Which uses of AI carry the risk?

The everyday ones, so that is where we train. When we built the classification behind our AI rules, we expected the hard cases to be the sophisticated ones: several agents chained together and connected to systems. Even a meeting summary turned out to be hard to judge, because a summary can hold a client's strategy or something under a confidentiality agreement, and the tool is the same one used for an internal stand-up.

So the literacy course every role goes through spends its time on everyday use. We took the rule through the leaders of each part of the company before we published it, so it arrived as something already discussed. The usage data shows the result: our security function has the highest share of the enterprise version of the tools anywhere in the company.

For a reader at this stage. Train for the ordinary use of AI, and take the rule through your leaders before you publish it.

We could have set up an AI department. Instead, we looked for the people who were already furthest ahead and gave them time to bring everyone else along, because a skill that lives in one team never reaches the projects.

— **Konrad Bałys**
Head of Engineering

The role widens as a result. An engineer keeps one deep specialism and adds the skill of building with AI across the steps around it, such as analysis, testing, documentation and review, which leads to a smaller, T-shaped team.

The skills programme

We built the programme as a system, with an expectation, a supply of learning and an owner for every delivery role.

How do we define AI skills? We no longer equate them simply with familiarity with specific tools. According to our definition, **AI skills encompass the conscious, everyday use of AI, an understanding of how AI can be applied within one's own area of specialisation, and the ability to evaluate the quality of AI-generated output**. They include also experimenting with new ways of working, sharing what works, and spotting what can be automated.

The AI competencies expected in each delivery role are described in our skills portal, **where people assess themselves against them**, and working with AI is written into the responsibilities and requirements of each position.

The expectations went into the annual goals of every competence centre and department, so the priority sat where people are measured.

At an all-hands meeting we asked everyone in a delivery role to assess themselves, and explained why: this is now the standard for how we deliver. From May, delivery leaders carried the expectation into their teams, first by working with AI themselves and then by asking their people to build the skills and use them in daily work. The formal standard covers delivery for now. Outside delivery, teams are invited to experiment, as in stage two.

The skills portal builds the expectation in layers. Every role shares a universal base of three skills:

- AI and LLM fundamentals,
- prompt engineering,
- AI security and trust.

On top of it come a practical skill, the main skills specific to the role, and additional skills that depend on the project.

People assess each skill on four levels: None, Newbie, Practitioner and Master, by how they use it in their daily work rather than by how much theory they know. The test for Practitioner is one question*: if I joined a project tomorrow that needed this skill, could I use it independently, responsibly and regularly?*

ROLE FAMILY	PRACTICAL SKILL	MAIN SKILLS FOR THE ROLE
Developer / Engineer	AI-assisted development	AI coding tool proficiency
Quality engineer / QA specialist	AI-assisted development	AI quality engineering
PSE / Architect / Tech lead	AI-assisted development	AI-assisted SDLC; spec-driven development; AI coding tool proficiency
Business analyst / UX	AI-assisted analysis and delivery work	AI-based knowledge and product context management
Consultant	AI-assisted analysis and delivery work	AI business value assessment; AI communication and stakeholder transparency
Project manager	AI-assisted analysis and delivery work	AI-enabled project delivery; AI communication and stakeholder transparency
Team leader	AI-assisted analysis and delivery work	AI-enabled team leadership; AI communication and stakeholder transparency
Delivery manager	AI-assisted analysis and delivery work	AI delivery governance and maturity management

AI skills recommended for each delivery role in our skills portal. Every role also shares the universal base, and additional skills depend on the project.

> *In delivery, working with AI is now a standard in every role and our leaders lead by example. We invite people into this change, and we expect it. But they have to want to be in it.*

— **Marta Gołowkow**
HR Business Partner

How it works in practice. Before anyone in the company gets access to the tools, they confirm they have completed an AI literacy course. It covers AI fundamentals, risk classification under the AI Act, our own policy and usage rules, and threats and safety.

Engineers in delivery also follow a modular self-study programme on AI-assisted software development. Each package pairs concepts with hands-on exercises in a set order, taking people from using an AI assistant to building with agents on our platform. The first modules are general enough to help in other roles too. Delivery also runs training for selected groups.

Training is only part of it. When a new way of working is ready, we introduce it in a team's current project, as the standard way of doing a specific task. People first see it working on their own project, and only then do we ask them to use it themselves.

Who owns it. The head of each competence centre is responsible for the AI skills of the people in that centre, and the centre's progress is scored every quarter. The AI Masters and AI Champions carry the skills into the projects. Nobody in delivery was excused from learning: the expectation went into the same review as everything else people are measured on.

AI Coffee: the community that was there first

Much of what the skills programme later spread was already happening in the company. In April 2024, almost a year before the programme began, Ewelina Magrel, our AI Practice Manager, started AI Coffee.

These are short, open meetings, each on a specific AI topic. They are run by a colleague who has tested a solution in their own work. First they show how they use it, then they answer questions and discuss it with the audience.

Every session is recorded, and the recordings, materials and links are kept in a shared channel that anyone in the company can join.

The topics grew with the company: from the basics, such as prompting, security and coding with an AI assistant, to experience from client projects, the requirements of the AI Act and working with AI agents. Guests came from outside the company too, Microsoft among them.

By the time the transformation began, AI Coffee already had a regular audience and a fixed place in people's calendars. The programme could use it to reach people outside formal training. New tools, rules and methods were presented to the whole company by people who already used them in practice.

The sessions also drew people from outside delivery. In March 2026, more than 160 people joined a session on building a working application without writing code, simply by talking to AI.

More than thirty recordings are now in the channel, and they remain available to anyone who joins the company later.

AI Coffee is one of three ways we build AI skills across the company. The literacy course gives everyone the basics, the AI Masters and AI Champions bring new methods into projects, and AI Coffee keeps the whole company talking about AI.

WORTH COPYING: PEOPLE

An open session on one topic

The move. A short, recurring meeting on a single AI topic, open to the whole company, presented by a colleague who has tried it and followed by questions. Every session is recorded and kept in one channel.

How you would use it. Start it before the formal programme, with the people who already experiment. Keep the format fixed and the bar for presenting low, and record every session, so the archive also serves the people who join later.

Watch for. A programme that builds its own channel next to a community that already works. Give the community a place in the programme, and let it carry the programme's news.

> *We started AI Coffee so that anyone in the company could learn about AI from a colleague who had already tried it, in the time it takes to drink a coffee. By the time the transformation began, people already knew where to bring their questions.*

— **Ewelina Magrel**
AI Practice Manager

AI Masters and AI Champions: who brings AI into projects

The open community helped people learn. But for new methods to reach projects and actually be used there, we needed more. So we introduced two functions: AI Master and AI Champion.

We did not create new positions. Both functions are taken on by our own people alongside their main responsibilities, with no change of title or reporting line. That is why the whole system was up and running in weeks rather than months.

AI Masters: expert knowledge. Our AI Masters are experienced specialists from our own delivery, usually architects or people in similarly senior roles. They keep doing their main job and act as AI Masters for an agreed number of hours each month, funded by the company.

At first, AI Masters worked mainly as consultants: any project could ask them for help in introducing AI. They advised teams, ran workshops and reviews, and built solutions that teams could put to work in their projects straight away.

> *Once people became more knowledgeable, the questions to our AI Masters began to dry up, and that was exactly the result we wanted. The Masters could spend their time on improving the methods our teams use every day.*

— Anna Galonska
Internal Project Manager who coordinates the AI transformation

Today, the main task of our roughly 35 AI Masters in the AI-Boosted Delivery programme is developing the company's standards for AI-assisted delivery.

Alongside the formal programme, the AI Masters have their own forum, Konsylium AI. At each session, one of them presents a working proof of concept to the rest of the company, together with the code, so that others can use it in their own work straight away.

AI Champions: support in every project. Every project team must have an AI Champion, usually its technical lead. The Champion makes sure AI is actually used in the project, helps coordinate that work and looks for the most effective ways to use it.

In practice, the Champion applies the standards developed by the AI Masters. They run and document experiments, support their teammates and make sure work with AI follows our policies.

Together, the two functions also form a natural path for growth. Everyone starts by using AI well in their own work. Those who want to go further can become the AI Champion in their project and, with enough experience, join the AI Masters and help shape the company's standards.

WORTH COPYING: PEOPLE

Experts for the whole company, a champion in every project

The move. Senior practitioners spend part of their time developing the standard and answering questions. In every project, one person makes sure the standard is used there. Everyone kept their main role and took this on as an additional function.

How you would use it. Pick the people who already have the skills and set aside part of their time for sharing what they know. Name one person in every project who is responsible for bringing AI into it. Give both groups one shared channel, so they don't give conflicting advice.

Watch for. Requests for help drying up. Ours did once people became more knowledgeable, and it was the signal to move the experts on to developing the methods in the use-case library.

Experiments become part of the goals

In stage two people were invited to experiment. In this stage experimenting became part of the goals of every team, and every experiment went into the register.

A failed experiment is as valid an entry as one that worked, and every entry records how its effect was measured and what came of it.

> **WORTH COPYING: AI GOVERNANCE**
>
> Experiments you write into goals
>
> **The move**. Experiments written into team goals and logged in an open register that takes failures as well as successes, each with what was measured.
>
> **How you would use it**. Ask every team, delivery or not, to try AI on one task it already does and record what happened. Keep the register open to every entry, failed ones included, and ask each one for a simple measure of what happened. Ask for a validated result only when someone proposes it as a use case.
>
> **Watch for**. A register with a high bar for entries. People log only the wins, so you learn nothing about the attempts that failed, or about who tried at all.

The Masters read the register, and an experiment that repeats across teams with good results is the raw material a use case is written from (stage four).

The register now holds 266 experiments from every part of the company: HR, recruitment, sourcing, marketing, IT and finance among them, as well as delivery. We read that number as evidence of openness and of movement from the bottom up.

Non-technical teams built working things.

> *What surprised me most about working with AI was how much one person can now do. Things that took me hours earlier in my career I can now do in minutes, at comparable quality, as long as I know the work well enough to judge the result.*

— Konrad Bałys
Head of Engineering

How we measure the skills: the AI Maturity Assessment

A survey tells you what people say. To follow the skills upgrade itself, we score it regularly with the people who run the work. The AI Maturity Assessment started with a baseline reading and has been repeated every quarter since. Unlike the survey, which covers the whole company, the assessment covers delivery only: its competence centres and client projects.

The head of each competence centre scores the centre in four areas:

- People: how the centre builds and spreads AI skills among its members;
- Practices and methods: how far AI-supported ways of working are defined for each role;
- Tooling, data and compliance: the guidance and assets that make AI use safe;
- Governance: how deliberately AI adoption is managed.

Each area is scored on five levels: Exploratory, Applied, Standardised, Strategic and Transformational. Delivery managers use the same method to score each client account they run. This case study reports the competence-centre view.

Each level is defined by what can be observed, so that two people scoring the same centre reach the same answer. For example, level three (Standardised) in the people area means that:

- the competence centre has defined the baseline AI skills expected for the role;
- introductory training, guidelines or curated learning paths exist, and members know about them;
- good practices are collected and shared across the community;
- new members receive AI onboarding relevant to their role.

WHERE OUR SCALE COMES FROM

We did not build the assessment from scratch. Delivery reviewed several published AI adoption and maturity models and chose the Google AI Adoption Framework as its reference model, because it fitted the way we deliver software best.

We adapted it to our needs. The original describes three phases of maturity: Tactical, Strategic and Transformational. We extended them to five levels, so that a single competence centre or client project can be scored precisely, and adjusted the areas to the way our teams work.

The framework focuses on technical and process capabilities. That is why it underpins the maturity assessment, while the four stages of our transformation come from our own model.

The skills upgrade in the competence centres, by area

AI Maturity Assessment, competence-centre view, three quarterly readings. Area averages across nine centres, each scored by the centre's head on a scale of 1 to 5.

How to read the scale

1 · **Exploratory:** isolated, uncoordinated experiments, no shared practices

2 · **Applied:** local initiatives, but no shared standards

3 · **Standardised:** repeatable practices and clear rules

4 · **Strategic:** AI built into the work across roles and aligned with strategy

5 · **Transformational:** AI as the default way of working

Three of the four areas rose in every reading. People rose from 2.56 to 2.89, governance from 2.44 to 2.89, and practices and methods, the weakest area in every reading, from 2.11 to 2.56.

Tooling, data and compliance was the strongest area in the first two readings and came in lower in the third, after teams had worked through the area in depth. The assessment's own report put it this way: "*the more we do and learn, the more clearly we see what still needs to be improved.*"

The assessment is scored by the people who run the work, which makes it a management conversation as much as a measurement. The year's actions are tied directly to what should move by the end of it.

In the words of Konrad Bałys, our Head of Engineering, *it is hygiene: it does not make the work good, but without it nobody can say where the work stands*.

KEY TAKEAWAYS FROM STAGE THREE

- **Make AI a skill inside every role.** A separate AI team is not enough, because a skill that lives in one team never reaches the projects.
- **Build the skills programme as a system**: an expectation per role written into goals, a supply of learning, and an owner in every department.
- Give the people who already have the skill time to teach it, and name one owner of adoption in every team, as a function within their current role.
- **Write experiments into team goals** and let the register take failures. The ones that repeat with good results have the potential to become your standard.
- **Score the skills every quarter**, with the people who run the work, against a rubric that says what each level looks like.
- Build on the places where people already share what they learn about AI.

Stage four

Organisational resilience

When does AI actually change how the work gets done?

Once the skills were spreading and AI-native delivery had become the standard way our delivery teams work, we started asking a different question: **how do we make the new way of working survive the next changes in AI?**

As long as a gain sits in one person's tooling, nothing about a project can be repeated. A company whose methods depend on one expert or one vendor is exposed to every change.

AI really changes how the work gets done only when a method stops belonging to one person, and the company can keep developing it without losing control. That is what resilience means: the delivery model, the tools, the governance and the commercial model can all absorb the next change in AI.

The board had already set the direction. We would earn more from the value we deliver, the components we reuse and the results we commit to, and less from hours worked.

The delivery standard reached its first version, our agentic platform followed, and the number of delivery methods approved for reuse grew from six to eighteen.

How does an experiment become a company standard?

One successful experiment in one project is only a start. Before other teams can use the method, someone has to describe what it delivers, what risks it carries and under what conditions it works. Someone experienced also has to review it before it reaches other projects.

Without a person looking after this process, good ideas stay with the people who had them. So we appointed an owner for the route and split it into three steps, which the experiments from stages two and three pass through. The requirements rise at each step.

- **From idea to experiment.** A team tries the idea in a project and logs it in the experiment register, with what it measured and the result. Every experiment goes in, including the ones that failed.
- **From experiment to use case.** The AI Masters review it, the author writes it up to a common template, and its effects, risks and conditions for reuse have to hold in more than one project. Use cases go into the library, open to every delivery team.
- **From use case to delivery standard.** The use case is built into everyday practice and onboarding, so that any team can request it, apply it and measure its effects without needing the author.

The bar starts low, because any attempt can go into the register, and is set high only at the door of the library. That keeps the register open to everyone and makes the library trustworthy.

Two kinds of standard meet here. The delivery standard is the way our delivery teams run a project: AI-native delivery, described later in the report. A use case is a single method that has passed the bar for reuse, and the library keeps growing inside the delivery standard.

From an idea to the standard

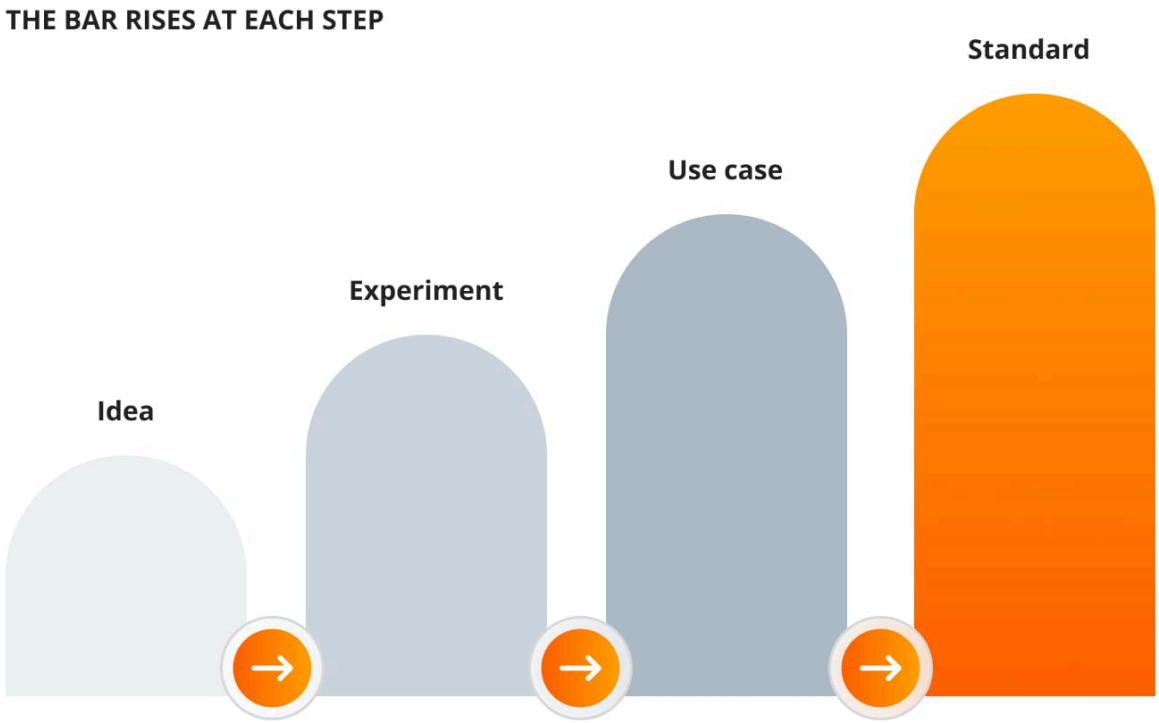

Idea	Experiment	Use case	Standard
Anyone, in any team	Tried in a project and logged	Reviewed by the Masters, written to the template, validated in more than one project	Promoted into practice and onboarding
No bar: an idea costs nothing to have	No result required; failures count; 266 in the register	Effects, risks and conditions for reuse stated; 18 approved, up from 6	Any team can request it, apply it and measure it, without meeting its author

Every tool has to be replaceable

Vendors may reprice at any time, and the running cost of AI is set to outweigh the cost of building with it, as happened with cloud.

So we took a multi-vendor position: every tool we use has to be replaceable, so that a vendor's outage or repricing has as little effect on delivery as possible.

The position asks something of us. We maintain several tools side by side, keep each one configured and compliant, and keep the shared standard of work independent of any single tool.

WORTH COPYING: AI GOVERNANCE

Choosing your infrastructure

The move. Four options weighed in the plan and put to the board: the productivity-suite cloud we already ran, a hybrid with a small local sandbox, a hardware partnership beside the cloud base, and full on-premises, each scored on pace of rollout, security posture, what clients in regulated industries would accept and which engineers would want to work here.

How you would use it. Decide it as a board, against what your clients in regulated industries will accept and who you want to hire, and not only against the cost column.

Watch for. The cheapest option. It is also the one that makes you look like everybody else facing the same decision, and that cost does not appear in the cost column.

Clients start asking for the method

As the method matured, clients began asking to learn our methods. The rule from stage two holds: the contract sets who owns what, both what the project produces for the client and our way of delivering, the library and the platform, which stay ours.

The method has become a competitive advantage and part of our intellectual property. It is what distinguishes our delivery from a supplier who has handed out assistants and stopped there, and we protect it accordingly.

The same redesign across the whole client process

The redesign did not stop at delivery. We rebuilt the steps that come before a project starts in the same way, so that the whole path of working with a client changed: answering the first enquiry, preparing the proposal, assessing the contract and putting the team together.

Lead qualification

More than half of the enquiries reaching the sales team were newsletters, spam or messages with no intention of buying, and sorting them by hand delayed the answers to the ones that mattered: the first contact took 20 working hours on average.

An AI system now reads each incoming message and judges its intent, builds a profile of the company behind it, registers qualified leads in the CRM with the right salesperson, and hands that person a package of context, including the case studies most relevant to the client.

90%

of enquiries sorted automatically

-60%

leads that could not be reached

Source: lead qualification case study

About 90% of enquiries are now sorted automatically, first contact takes 6 working hours instead of 20, and the number of leads that could not be reached fell by 60% (the lead qualification case study).

Proposal preparation

Demand for quotes was growing faster than a manual quoting process could handle, and the presales team spent much of its time on repetitive work.

The team now works with a set of AI agents built in-house.

- one analyses each incoming enquiry and creates the CRM record;
- another studies the client's documents, searches past proposals and drafts standard content;
- another generates the structure of the project, from functional areas and the backlog to effort estimates, risks and the suggested team.

-56%

time to prepare a full proposal

+40%

work handled by the same nine-person presales team

Source: AI presales case study

Preparing a full proposal now takes 56% less time, and the nine-person presales team handles 40% more work (the AI presales case study).

Contract risk assessment

Assessing the risk in a client contract was manual work that took days, and there was no common measure for comparing the risk of one contract with another.

An AI system now reads the contract, compares its clauses with a central risk register kept by our legal, delivery and finance teams, classifies the risky clauses and scores the contract's overall risk profile.

The preliminary assessment now takes minutes instead of days, almost 90% less time, and every contract can be compared on the same risk score, which speeds up negotiation. The contract still goes to our legal team for the final review.

Matching people to projects

Finding people with the right skills and domain experience for a project meant checking several separate sources: employee CVs, the contractor base and the list of people between projects. An

AI agent now searches all of them in one place and helps resource managers judge who fits a project. It is used every day in staffing decisions.

IF YOU ARE STARTING OUT

Look for early wins in the processes either side of the core work, such as intake, quoting, contracting and onboarding. They are internal, they repeat many times a week, and their effect is easy to measure.

THE SECURITY QUESTION AT THIS STAGE

How do you keep the rule current when the tools change every month, and easy enough that people use it?

By assessing the use as well as the tool, and building a tool that applies the rule.

The rewritten AI policy shipped with a Compliance Assistant: anyone describes what they want to do with AI, on what data and for whom, and gets an answer in a conversation. A fixed classification sits underneath, with around 150 everyday uses assessed in advance, so that common cases get consistent answers. Uses the classification does not cover are assessed separately.

Low risk means go ahead. Medium risk is logged, so the company knows what it carries. High risk goes to the level that holds the mandate to accept it, and usually comes back with a way to lower it. Accepting a risk internally never replaces what a client's contract or the law requires.

Around it run configuration guardrails on the centrally provided tools, monitoring, ongoing education, and tools that can be swapped, to limit the effect of any one vendor's outage or repricing on delivery. Where somebody builds an agent that acts with their own permissions, a person is placed where the output can be checked.

“The goal of our AI rules was never for people to use AI less. We wanted them to use it more widely, without doubting whether they may.” Dominik Samociuk, Head of Security & Compliance

For a reader at this stage. Use the technology to make the rule easy to follow, then watch whether people follow it.

The commercial model changes with the process

A delivery process redesigned around AI changes what a client is paying for. When a team of three delivers what used to take ten, an hourly rate measures the wrong thing. **So we tie new work to business outcomes:** what the client wants to achieve, and the success measures agreed with them at the start.

Outcomes depend on understanding the client's business as well as the technology. Our consultants and engineers bring deep knowledge of the domains we work in most, among them insurance, finance, media, and energy and utilities.

The outcomes we commit to are the ones that matter to the client. AI makes delivery faster. Domain knowledge decides what is worth delivering, and how best to do it.

To become resilient, a company has to redesign its key processes and, at the same time, work out what that redesign does to its business model.

> *"When we take responsibility for the result, we decide how the work gets done, and that is when our own components and AI solutions can be used to the full. Step by step, our business model will depend less on the number of hours worked and more on what we have built, invested in and can use again."*

— Paweł Pustelnik
Vice-President of the Board and COO

KEY TAKEAWAYS FROM STAGE FOUR

- Appoint someone to turn successful AI experiments into a shared standard of work. Without them, good ideas stay with the people who had them.
- Collect every AI experiment in one place, failures included. A simple measure is enough to start with. Ask for hard evidence only from those meant to become a standard.
- Don't depend on a single AI vendor. Choose tools you can replace, and set up the way your teams work so it holds up whichever tool they use.
- Make your AI rules easy to follow before you ask people to follow them. Send decisions about risky uses to the people with the mandate to make them.
- When you redesign your key processes around AI, redesign your business model too. Once the same work takes fewer hours, billing it by the hour stops making sense.

AI-native delivery

How we rebuilt software delivery around AI

Software delivery is one of the core parts of our business, so this is where the transformation had to prove itself. **AI-native delivery is now the standard way our delivery teams work.** AI is built into every step of a project: discovery, design, development, review, security and governance.

Our traditional delivery process is still available. At the start of each project we explain to the client what the AI-native flow would change for them, and together we agree which of the two flows the project will follow. A client may prefer the established way of working, or their contract may restrict the use of AI.

What matters most in redesigning the process of software delivery is how deeply AI is built in. When AI is added on top of an old workflow, single tasks get faster and the process stays the same. When AI is built into the workflow, the whole project changes: the specification, the team, the review and the pace.

Everything starts with the specification. The build now follows the specification closely and quickly, so the specification matters more than before. We agree the goals, the success measures and the requirements with the client at the start. Getting them right takes the knowledge and experience of people on both sides.

> *More and more clients come to us with a working proof of concept of their own. It is the new way of saying what you want. Somebody still has to take it to production safely, and that is what they now ask us about.*

— **Daria Polończyk**
Head of Delivery

The framework: crawl, walk, run

Our framework builds AI into every step of the work. It has three levels, and each level describes how much of the work AI does and how much a person does.

Crawl: AI supports people. AI drafts, summarises and checks. For example, it turns notes from discovery workshops into a first version of the documentation. People write and decide everything that matters.

Walk: AI is part of the team's standard way of working. Requirements are written as a specification that AI can build from. Every change is checked automatically before a person reviews it. The team measures what AI contributes.

Run: AI does most of the building. Agents build a feature from an agreed specification, inside fixed guardrails and automated quality gates. A person reviews and approves the result instead of writing it.

> **WORTH COPYING: PROCESSES**
>
> **Crawl, walk, run**
>
> **The move**. Three levels of how far AI goes on a project: AI supports people (crawl), AI is part of the team's standard way of working (walk), or AI builds and people review (run). The level is set at the start of each project.
>
> **How you would use it**. Decide the level for each project or workflow at its start. Base it on what the contract allows, how sensitive the data is and what kind of code is involved. Write the decision down, so that the team and the client know it before the work begins.
>
> **Watch for.** Treating run as the target for every project. The right level is the one the project's conditions allow.

The level is set for each project at its start. It depends on what the client's contract allows, how sensitive the data is and what kind of code is involved. A project with sensitive data may stay at crawl from start to finish, and that is the right level for it. The same team can work at run on one project and at crawl on the next.

Within each project, an oversight rule applies to every piece of code: the closer the code is to the heart of the client's business, the more of it a human writes.

Two ways to run a project

A **gated lifecycle** moves through fixed phases. AI supports every phase, and a person signs off at each quality gate before the project moves on.

A **rapid loop** builds, tests and adjusts in short cycles, and repeats them until the solution does what the client actually needs.

We choose the approach with the client at the start. Most projects combine the two: they move fast where that is safe and follow the gates where precision matters.

The platform behind it

Both approaches run on our own agentic platform, which has two parts.

The **engine** works in two phases. First, people and AI write the specification together. Then agents build the solution from it. A person oversees the work throughout, and nothing is released without a person's sign-off.

The **library** holds the skills and agents that all our developers share. Every team works to the same standards and safeguards, so the quality of a project does not depend on who happens to be on the team.

The team shape

Teams are smaller. A traditional team had six to ten specialists. Each did one part of the job and passed it on to the next, and production itself was the slowest and most expensive step.

An AI-native team has two or three people. Each of them brings one deep specialism from a competence centre and the skill of building with AI agents on our platform. One person can now cover ground that used to need several hand-offs.

Discovery and building overlap. Prototyping starts while discovery is still under way. **Production stops being the bottleneck, and coordination takes its place.**

Experience counts as much as ever. Agreeing goals and success measures at the start is where twenty-five years of delivery experience makes the difference. What has changed is how the project is carried out after that.

The client sets the pace. A small team that builds to a specification mostly waits for two things: decisions and acceptance. So before the project starts, we agree who on the client's side can decide within a day and who accepts each increment. A regulated company can arrange both without loosening any of its controls.

Quality without compromise

In traditional delivery, requirements are scattered across emails, documents and calls, and the backlog is built and estimated by hand. Code review happens only once a piece of work is finished, and the scope of testing depends on how much time and budget the project allows.

Now, AI agents gather all of that input into one structured body of project knowledge. AI helps to build and size the backlog. AI review runs at every stage, and automated tests cover units, contracts, APIs and complete user journeys.

Every speed-up is checked. When one step gets faster, we look for the step that now holds the project back, and we fix that one too. Code review is the obvious example: if AI writes ten times more code and people still read every line, review becomes the bottleneck.

Quality stays where it was. Our quality goals were high long before AI arrived, and they are now built into the tools. Keeping that level while the amount of code grows is the real challenge.

> *AI-native delivery is not random vibe coding. We get the pieces right first: what is to be done and how it is designed, and only then deliver it at the level of quality we have always held.*

— Konrad Bałys
Head of Engineering

What the method has done on real work

These are a few examples of methods from our use-case library, which holds eighteen. Each has been used on more than one client project, its effect has been measured every time, and any delivery team can apply it.

METHOD	BEFORE	AFTER
Technical documentation generated from the code of a multi-module system, checked by a second agent before a person reviews it	2 to 3 days per subproject	Under 30 minutes
A migration skill that upgrades a data platform and produces a change a person can review, used on around 50 data products	About 2 days per data product	4 to 6 hours
End-to-end development from ticket to reviewed change, with the developer approving key decisions and the final change	Manual rewriting of tickets and multi-stage implementation	About half the time
Discovery run as a Git repository, with AI drafting the analyst's documents from meeting notes and client materials.	120 hours	About 55 hours

From the use-case library. Every figure is signed off by the method's owner.

Two client projects show what the method does on paid work, in two very different situations: a fixed-price system whose requirements kept changing, and a new system with a scope fixed from the start.

International strategy consulting firm: a fixed-price web system

The challenge: The client wanted to automate its HR processes. The system had to serve several user groups and integrate with existing tools, and its requirements evolved throughout delivery as the business refined its needs. Price and a proven ability to deliver were the deciding factors in choosing a partner.

What we did: We agreed a fixed-price engagement and delivered with its specification-driven AI-native delivery model. This was the first client project to use the whole framework. This let the team deliver the full scope quickly, including every requirement that changed along the way, while maintaining its established engineering and quality standards.

The result: The client received the full agreed scope on time, to a high standard of quality and within the fixed price. Delivery took more than 30% less effort than Future Processing's historical benchmark for a system of that size. That efficiency is what made a competitive fixed price possible. The client then extended the engagement.

What we learned. When engineering speeds up, everything around it has to speed up too: business decisions, ready requirements and acceptance testing. Otherwise the project waits for them.

Fortaco: a working prototype in two weeks

The challenge. Fortaco needed a blockchain-based solution, integrated with SAP S/4HANA, to record certified-origin steel received into inventory. The scope and the budget were fixed from the start.

What we did. We wrote the full specification first. AI agents then built the application from it, and a person reviewed and approved every step.

The result. Within two weeks the client had a fully working prototype, on which the first tests of the integration with SAP could begin, and working on it from that point changed the way we work together. The client then received a finished, accepted system within the fixed budget, with a twelve-month warranty (the Fortaco case study).

Why it worked so well. The project suited the AI-native delivery method: a new system, a closed scope and a well-known domain.

How do we keep client code and data safe when AI is part of the work?

Our rules for using AI on client code and data always start from the client's contract. They are the same on every project. They are written down, and the client's own reviewer can read them in a minute.

Oversight depends on the kind of code. For core business logic, such as payments, AI only advises and people write every line. For other business logic, AI writes the code and a person reviews it. For routine code, AI does the work and a person gives the final approval. Quality checks run on every commit, whatever the kind of code.

Agents are supervised too. When a team builds an agent that acts with a person's own permissions, a person checks its output before it counts.

We assess the use of AI as well as the tool. New uses are checked with our Compliance Assistant, and uses outside its classification are assessed separately. The risk depends on the tool, how sensitive the data is, how it is processed and what the contract says. A high risk usually comes back with a way to reduce it, such as a different tool, anonymised data or a clarified clause.

> *We assess AI risk through both the tool and its intended use: the data involved, the contractual context, the safeguards and who is accountable.*

— **Dominik Samociuk**
Head of Security & Compliance

Client data stays the client's. The use of client code and data with AI stays within the agreed contractual scope and safeguards. Training and grounding are assessed separately, and approval for one does not imply approval for the other. On the centrally provided tools, configuration guardrails apply parts of the policy in the tool itself, as well as on paper.

The client's rules come first. Where a client's AI rules are stricter than ours, theirs apply on their project, and AI is used only within what the client has agreed. Many legal teams are now writing their first AI clauses, so we expect this conversation and come prepared for it.

Regulation shaped the programme from the start. The AI Act, GDPR, NIS2 and DORA were considered in its first plan, each where it applies to us or to a client's project.

KEY TAKEAWAYS FROM AI-NATIVE DELIVERY

- **Rebuild the way your IT teams work around AI** rather than adding AI to the old process. Only then do the specification, the team, the review and the pace change together.
- **Invest in the specification**. When AI builds quickly and closely from it, its quality decides the result.
- **Agree at the start of every project how much of the work AI will do**. Consider how sensitive the data is and what contracts and regulation allow. Keep the traditional way where it is needed.
- **Write down your rules for using AI on code and data**, briefly and clearly. Client and regulatory requirements come before internal rules.
- Before the project starts, **agree who on the business side decides** within a day and accepts each part of the solution. A small AI-native team mostly waits for decisions.
- **Measure the effect of every method before you roll it out** to other teams.

Part three

Where to start, and what to do next

The playbook

The table describes the four stages of AI transformation. Find the one that best fits your company today and start with its recommendations. Once the next stage describes your company better, move on to its recommendations.

STAGE	YOU ARE HERE IF	WHAT WE RECOMMEND
1 · Point adoption	Some people use AI on their own, but nothing is shared or measured. The risk of shadow AI, meaning AI tools used without the company's knowledge, is growing.	1. Decide at board level what AI should change in your company. 2. Measure how people across the company use AI, and plan a second measurement. 3. Plan with the people who will do the work, and let departments move at different speeds. 4. Set goals and a budget, and review the plan every quarter.
2 · Wide adoption	The board has decided that AI should cover the whole company, but there is no shared way to get tools yet, and no rules for using AI.	1. Give every team a fast, funded way to get approved AI tools. 2. Publish a short AI policy that says who is accountable for AI-assisted work, with a review date. 3. Track usage and spend from the first licence. 4. Invite every team to try AI on its own work.
3 · Shared AI competence	Everyone has approved tools, but a few people get far more out of them than the rest.	1. Make working with AI part of every role. 2. Write AI skill expectations into each role's goals, with training to match. 3. Give your strongest AI users time to teach, and name an AI champion in every team. 4. Write experiments into team goals. 5. Assess AI maturity every quarter against a fixed rubric.
4 · Organisational resilience	Teams work with AI in a shared way, but it rests on a few people or a single vendor.	1. Name an owner for turning successful experiments into company standards. 2. Make sure you can replace any AI tool or vendor. 3. Keep the AI rules simple to follow, and decide each risk at the level that owns it. 4. Redesign your key processes around the way people now work with AI.

What comes next: from skills to ROI

We set out to raise our people's AI skills, and the measurements show that progress. The business effects are already visible too: in single tasks, in individual projects and in our processes. What we cannot show yet is one company-wide figure for return on investment. We will publish it once the new way of working covers most of our projects and has run for a full year.

What is measured, and when

Measured	Operational level	How long a single task takes
	Tactical level	How much faster a team delivers a whole piece of work
Next	Strategic level	Company-level costs and results, once the new way of working has settled in

Effects on single tasks and on team delivery are measured now. Company-level costs and results become readable after a full year of working the new way.

> *Our first priority was that everyone should learn to work with AI and start using it. The next step was standardisation, so that teams would work with AI in a consistent way, turn use cases into real business value and scale it across client projects.*

— **Konrad Bałys**
Head of Engineering

What should you ask an AI delivery partner?

Ask to see the method in use. A good partner can show you a method it applies across projects, with its measured effects, risks and limitations written down. Anyone can show you one good project. Ask what the team does every day.

> *You will talk to one firm and hear that they have an AI assistant for everyone, and that is it. You will talk to another and hear that they have their own delivery lifecycle, a whole system, methods, security gates, an agentic platform that works around the clock, and experts who know how to keep it under control. That is what you are choosing between.*

— Daria Polończyk
Head of Delivery

At what scale does that run? A 96% saving on a task done twice a month is a number, but it has little effect on a project. Ask for the frequency alongside every percentage.

Which delivery model, and how much AI, will my project use?

A partner that has rebuilt its delivery can run your project the traditional way or in an AI-native model. It should also explain how much of the work AI will take on, and why. The answer must follow from your contract, from how sensitive your data is, and from the kind of code involved.

How will AI work with my code, and who decided that?

Find out which tasks AI will perform on your code, what it will have access to, and who will check its output. Ask how much of the work AI will take on and how much will stay with the team. You need to know how the partner assesses the risk of each use of AI in your project, and who approved it. Also find out what happens when the contract doesn't cover something. A list of approved tools is only part of the answer.

Who owns what AI produces on my project? Ask who will own the code, the data and anything the partner's tools generate. An experienced partner sets this out clearly in the project contract from the start: what belongs to you, and what stays with the partner, such as its platform and methods.

Who on my side needs to decide and accept, and how fast? In an AI-native project, decisions and acceptance set the pace. Agree before the project starts who can decide within a day and who accepts each increment.

What happens to the price as the work gets faster? Ask what the partner charges for: hours, or results. A partner who still sells only hours has every reason to keep the hours.

What does an AI-mature organisation look like in practice?

Fifteen months ago, we were working out how to take AI from a handful of people into the work of the whole company. Today the questions in our meetings are different: which proven method fits the task in front of us, how to make decisions faster now that the work takes less time, and how to keep quality where it was as output grows.

Each of the four stages moved us closer to that shift. Point adoption showed that the value sat with a few people. Wide adoption gave everyone the tools, the rules and an invitation to try. Shared AI competence made working with AI a standard for every delivery role, assessed every quarter. Organisational resilience made the methods the company's own: any tool can be replaced, and key processes have been redesigned around the way people work with AI.

> *AI will take over some of the tasks we do today, and those tasks are not all we do. The rest of our work will grow in value, and entirely new work will appear. Our job is to find that work, and to become the people who do it.*

— **Michał Sztanga**
CEO

About us

About Future Processing

Future Processing is a tech strategy advisor and tech delivery partner with 25+ years of experience. With our consulting mindset and deep domain expertise, we are obsessed with transforming your business ambitions into measurable results.

We proceed according to a unique, AI-enabled advisory and delivery framework grounded in our technological heritage and AI roots and continuously optimised to deliver faster. We design our solutions on the basis of clients' technology foundations and data infrastructure to high quality, governance and compliance standards so that they scale safely. We can modernise even complex legacy systems and operate in highly regulated industries.

More about Future Processing · Our AI and machine learning services

PLANNING YOUR OWN AI TRANSFORMATION?

Talk to the people who ran ours. We will look at where your company stands with AI today and what your next step could be. **Contact us →**

How we measured

A note on method

As of October 2026. This case study was published in October 2026, and all figures describe the situation at that point. The survey covers two waves (summer 2025 and April 2026), and the maturity assessment covers three quarterly readings from January 2026. AI tools and the ways of working with them change quickly, so later readings may show a different picture.

Two instruments, and what each reads. The employee survey reads adoption: what people do with AI and how often, alongside what they want from the company and what they feel they are given. The quarterly AI Maturity Assessment reads maturity: how far skills and practices have been built, scored against a rubric of observable indicators. This case study uses both, as well as the delivery registers.

The survey. Wave one ran in the summer of 2025 with around ninety questions across the whole company, 446 respondents and a 66% response rate. Wave two ran in April 2026 with a third of the questions and the comparable core kept: 480 respondents, of whom 293 had answered wave one, the matched panel. A third wave is planned. The first wave and the two-wave design were the work of Mike Zamarski, our AI Strategy Advisor, who also ran the workshop cycle; the second wave and the analysis were designed and run by Remigiusz Tunowski, our BI Consultant. The survey figures are self-reported.

The usage detail. Daily use across the whole company doubled between the waves, and rose in every delivery role by a factor between 1.7 and 2.5, with the order between the roles unchanged. The share of people placing themselves as enthusiasts or pioneers rose by about half, and the share placing themselves as sceptics fell by more than half. The self-reported hours saved are not used anywhere in the report, because a declared saving on tasks of unknown frequency tells a reader very little.

The maturity assessment. It runs quarterly, from a baseline in January 2026. The head of each competence centre scores the centre; delivery managers, with their team leaders, score each client account. Both use four areas and five levels, and each level is described by what can be observed. The levels adapt a published adoption framework, extended from three phases to five so that a single centre or account can be scored precisely. This case study publishes the competence-centre area averages. Account scores stay internal because they describe client work. Where the published summaries print a different overall average from the one their own area scores give, we use the area scores.

Why we measured ourselves rather than the market. Published figures on AI and productivity contradict each other, and a controlled trial does not work here, because early adopters select themselves. Measuring the same organisation twice, with the same instrument, gives a difference rather than an absolute, and the difference can be attributed to what happened in between.

The stages and the score. The four stages in Part two are our own description of how a company moves from AI adoption to AI transformation, built on the four-point AI Organisational Resilience Framework designed by Mike Zamarski. They are not a scale, and the case study never scores itself on them. The quarterly assessment is the instrument that scores.

Declared against measured. The survey figures are declared, and the delivery figures are measured in client projects and carry named owners. The report does not mix the two, and it prints no financial figure.

What fifteen months can establish. Fifteen months can establish that the skills programme reached every delivery role, that team-level practice moved further than anything else in the survey, that the competence centres' skills rose against a fixed rubric, and that specific methods now run in client work with measured effects. The effect on company-level operating and capital cost is the next reading, after a full year under the new way of working.