A PAPER FOR BOARDS

Start with the foundation

What boards and supervisory directors need to understand before their organisation builds AI agents

Wendy Morée and Lavanya Trehan
November 2026

www.techtogether.ai

Why we wrote this

Every board we speak with asks some version of the same question. The organisation is using AI, the budget line is growing, and sooner or later someone will propose agents. What does the board need to understand before it says yes?

We wrote this paper to answer that question in plain language. It applies whether you supervise a bank, an online retailer, a supermarket chain, a family owned funeral business or a company owned by private equity, because the question is the same everywhere. It is not a sales proposal. Our rates for 2026 are at the end for anyone who wants them, but the paper stands on its own. It sets out how we think about AI adoption, which questions we would ask if we sat in your seat, and where we believe the limits are.

One thing we want to say at the start. Nobody can tell you exactly when AI will pay for itself. The world is too uncertain for that, and right now the only thing we can be sure of is that it will stay uncertain. Anyone who promises you a return by a fixed date is guessing. What a board can ask for is evidence, gathered step by step, and the discipline to stop when the evidence is not there.

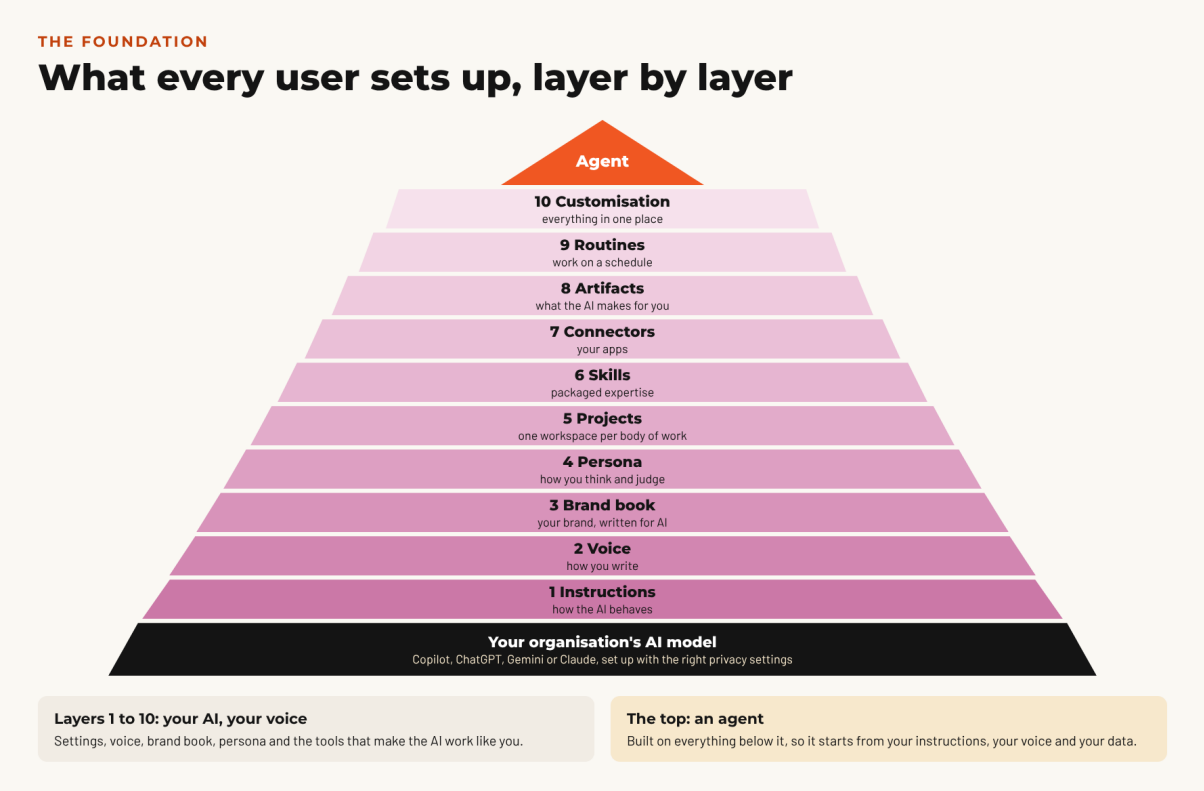
The foundation comes before the agent

Most conversations about AI start at the top, with the agent: the thing that answers customers, prepares meetings or scans contracts on its own. We start at the bottom. An agent is only as good as the layers it stands on, and in our experience an agent rarely fails because of the agent itself. It fails because nobody set up the layers underneath it.

We describe those layers as a pyramid. The foundation is the AI model your organisation has chosen, set up properly, with the privacy settings in place and the use of company data for training switched off where it should be. On top of that, layer by layer:

1. **Instructions:** how the AI should behave, what it should never do and how it should check its own work.
2. **Voice:** how you write, so the output does not read like a template.
3. **Brand book:** your brand written down in a form the AI reads every time it writes, designs or presents.
4. **Persona:** how you think and judge, so the AI makes choices you would recognise.
5. **Projects:** one workspace for each body of work, with its own documents and context.
6. **Skills:** expertise packaged once so it can be used again and again.
7. **Connectors:** links to the applications you already use, such as mail, calendar and CRM.
8. **Artifacts:** things the AI makes that you keep and share, such as a memo, a dashboard or an explainer.
9. **Routines:** the same piece of work, done again on a schedule, still checked by a person.
10. **Customisation:** everything above brought together in one place.

The agent sits at the top. It starts on its own trigger, prepares an action, waits for approval and keeps a log. It can only do that well because everything below it is in place.



It does not matter which model you use

Whether your organisation works with Microsoft Copilot, ChatGPT, Gemini or Claude, the layers are the same. The names of the buttons differ, the principle does not. That matters for a board, because the choice of model will change. Suppliers leapfrog each other every few months and prices move with them. A foundation that is built well, with instructions, voice, brand book and persona written down as company property, moves with you when you change supplier. A foundation that lives in one person's head or one vendor's settings does not.

So our advice is simple. Install the basics first, across the people who will use AI every day, before anyone builds anything more ambitious. It is less exciting than a demo, and it is where the results come from.

Why the board should go through it themselves

We recommend that executives and board members go through the same foundation training as their people, on their own laptop, with their own voice and their own work. Not so that they can build agents afterwards. An executive who spends evenings configuring agents is doing the wrong job. The reasons are different, and to us they are more important.

To negotiate with consultants and suppliers

Proposals for AI programmes are full of hours, licences and promises. A board member who has set up a model and watched an agent go from a job description to a working run knows what

takes an afternoon and what takes a quarter. That changes the conversation with every consultant and supplier at the table. It also helps to see through what Gartner calls agent washing: existing chatbots and assistants relabelled as agents. Gartner estimates that only about 130 of the thousands of vendors that claim to sell agentic AI offer the real thing.

To ask the right questions

A board does not need to know how to build. It needs to know what to ask. Who owns this agent? What did we measure before we started? Who checks the output, and how many hours does that cost? Where do the instructions and logs live, and are they ours? Those questions come naturally once you have done it yourself, and they are hard to ask convincingly when you have not.

To feel what adoption asks of people

There is a moment in every training when the AI confidently gets something wrong, and a moment when it saves an hour of work. Board members who have felt both understand why their people hesitate, why adoption takes longer than the slide said, and why the first weeks cost time before they give any back. That understanding shows in how the board supports the programme. McKinsey's 2025 survey found that organisations getting the most value from AI are about three times more likely to say their senior leaders visibly own and commit to their AI initiatives.

To decide whether you believe in the method

Last, a board has to decide which approach to AI adoption it wants to back. That is easier to judge from the inside than from a presentation. Go through it, then decide whether this is how you want your organisation to learn.

There is a legal reason too. The EU AI Act expects every organisation that uses AI to take measures to support the AI literacy of its people. It is hard to ask that of the organisation if the board has not done it first.

Train the supervisory board and the executive board together

Our advice is to put the supervisory board and the executive board in the same training, in the same room. They will be asking each other questions about AI for years to come, and those conversations go better when both sides have built the same foundation, made the same mistakes and seen the same results. It also takes the hierarchy out of the room for a day, because everyone starts at the same layer of the pyramid with their own laptop and their own work. And it gives the supervisory board a first hand view of what it will later be asked to supervise.

What the research shows

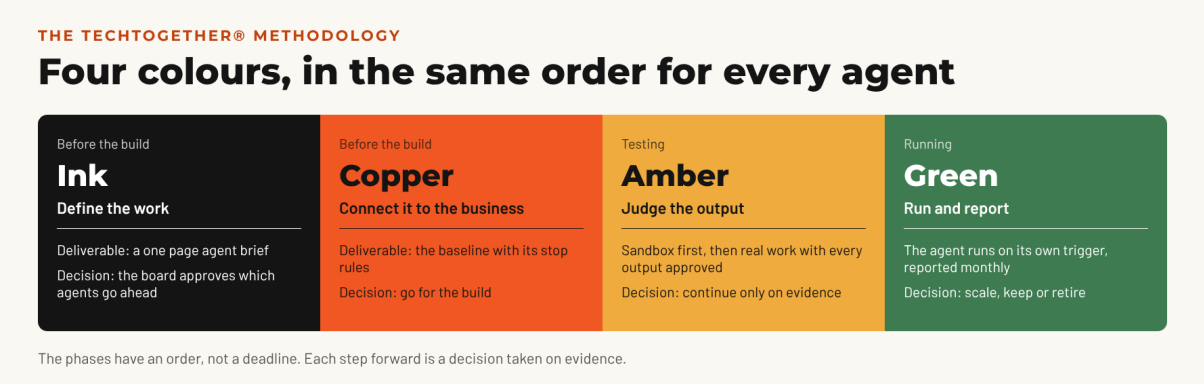
Most boards already have an advisor's view and their own reading, so we keep this short and only say what each source adds.

- **Wide use, thin results.** McKinsey's 2025 State of AI survey found that 88 percent of organisations use AI in at least one function, but only 39 percent report any effect on EBIT, and most of those say it is less than 5 percent. About 6 percent attribute 5 percent or more of EBIT to AI. Those few are around three times more likely to have redesigned their workflows, which is one of the strongest factors the survey found.
- **Results are possible.** Bain's 2025 Global Private Equity Report found that a majority of portfolio companies were testing or developing generative AI and that nearly one in five had use cases in operation with concrete results.
- **Many agent projects will stop.** Gartner expects more than 40 percent of agentic AI projects to be cancelled by the end of 2027, because of rising costs, unclear business value or weak risk controls. We read that as a governance warning more than a technology one.
- **The time is there to win back.** Salesforce's seventh State of Sales report, a survey of 4,050 sales professionals, found that they spend about 40 percent of their week selling and the rest on other work.
- **AI is uneven.** In a field experiment by Harvard Business School and BCG with 758 consultants, AI made people faster and better on tasks it handles well. On a task chosen to sit just outside what it does well, consultants using AI were 19 percentage points less likely to reach the correct answer. Tasks that look equally hard to a person can fall on either side of that line.

The TechTogether® Methodology: from foundation to a working agent

Once the foundation is in place, the question becomes which agents to build and how to keep them under control. We use four colours, in the same order for every agent. They are the same four colours as our training, so what an executive learns there is what the organisation does afterwards. Every phase has an owner, a deliverable and a decision, and the board sees the outcome of each one.

The phases have an order, not a deadline. How long each one takes depends on the data, the people and the workflow. Some agents move through in weeks, others need months. What matters is that every step forward is a decision taken on evidence, not on a date in a plan.



Ink: define the work

Before anything is built, the work is defined. The executive picks candidate workflows, ideally one close to revenue and one close to cost, and labels them that way. Each candidate is scored on five questions. How close is the output to a moment where money is won, renewed, lost or spent? How often does it run, given that an agent that runs weekly also learns weekly? Is the data it needs available and reasonably clean? Can the owner judge an output in a few minutes? And would a mistake that slips past the approval step be small and reversible? The highest scores go forward. For each, the person who runs the workflow today writes a job description for the agent. The deliverable is a one page agent brief. The decision is the board approving which agents go ahead.

Copper: connect it to the business

Every agent is then connected to the people, the data and the numbers around it. The owner, the builder, the reviewer and the person who keeps the data clean are named. The sources the agent will read are checked and cleaned where needed, which is often the largest piece of work and useful well beyond the agent. Finance records the baseline for the workflow, for example hours spent, response times, conversion at that step and cost per case. And the thresholds the agent has to meet at each gate are written down, together with the rules for stopping. The deliverable is the baseline with its stop rules. The decision is whether to build.

Amber: judge the output

The agent first runs in a sandbox on real but low risk material, and nothing reaches a customer. When that works, it moves to real work, and the owner approves every output before it is used. What is measured is how often an output is accepted without major changes, how long a review takes, whether the same mistake appears twice in the log, and how many hours are given back. At each gate the board continues only when the thresholds from Copper are met and the hours given back are greater than the hours spent reviewing.

Green: run and report

Only then does the agent work on its own trigger, with the owner reviewing a sample and anything the agent flags. The board then decides to scale it, keep it as it is or retire it. Retiring an agent that misses its criteria is the method working, because spending stops before customers or revenue are exposed. From then on, one page a month shows each agent with its stage, owner, leading signals, hours given back, running cost, incidents and the last decision taken. Together with a register of every agent, its purpose, data sources and approval step, that page becomes the evidence an auditor, a regulator or a future buyer will ask for.

WHO DOES WHAT

From the training to the report to shareholders

	After the training	Ink	Copper	Amber	Green	Ongoing
Shareholders	Expect quantified AI goals					Monthly page with leading signals
Board	Asks for a plan with stop rules	Approves the agents	Approves baseline and stop rules	**Gate decisions on evidence**	**Scale, keep or retire**	Agent register for audit and diligence
Executive	Sets the vision and the questions	Picks the workflows	Names owners and baseline	Reads the log every week	Monthly review with owners	Next agents
Agent owner		Writes the job description	Agrees the thresholds	Tests, then approves every output	Spot checks	Owns the result
Builder			Checks the data	Builds and tunes	Schedules the run	Maintains
Agent				Sandbox, then real work, all approved	Runs on its own trigger	Runs and learns

Return on investment, without a promise

Can AI agents grow revenue? Yes, under two conditions.

The first is time. The work changes quickly once an agent runs, but money follows the rhythm of each business. A company that sells to other businesses sees revenue move after one sales or renewal cycle. A company that sells through partners sees it when the partners sell more. A company that sells to consumers usually sees it first in customers who stay, because keeping a customer costs less than winning one. None of these rhythms can be forced by a better agent.

The second is honesty about the kind of value each agent creates. Some agents bring revenue closer, because they give sellers, partners or service staff time back for the work customers pay for. Others mainly lower cost. Both belong in the plan, in different columns. A board loses credibility with its shareholders the moment a cost saving is presented as growth.

That is why we usually suggest starting with one agent of each kind. The cost agent tends to show results earlier, and those results buy patience for the revenue side. In the meantime, the board watches the leading signals that appear long before the financial ones: hours given back, response times and the share of outputs approved without changes.

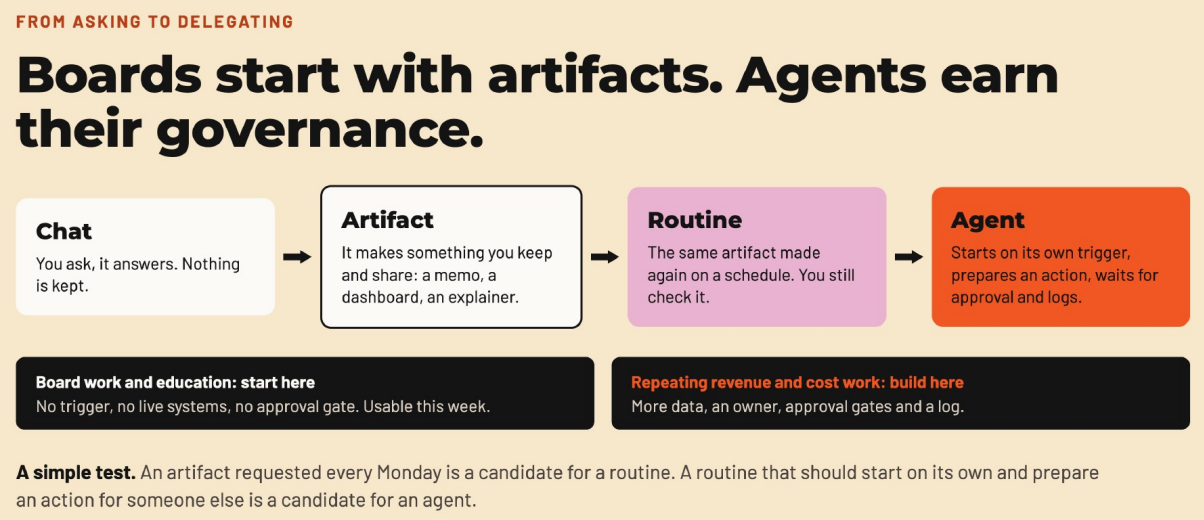
What this looks like differs per sector. For an online retailer it might be an agent that answers order questions in the company voice and hands anything sensitive to a person. For a supermarket chain, an agent that answers questions from stores about procedures and promotions. For a funeral business, an agent that prepares the paperwork after each arrangement so the staff can spend their time with the family. The method is the same each time.

Where artifacts are enough

Not every task needs an agent. For board work and for educating a leadership team, artifacts are the easier and safer place to start. An artifact is something the AI makes on request that you keep, edit and share: a summary of a long board pack with the open questions marked, an explainer that lets the board test the assumptions behind a forecast, a short brief on a

competitor before a strategy session. The person who asks for it also checks it, so it needs no trigger, no link to live systems and no approval gate. A board can use artifacts the same week.

A simple test tells you when to move on. An artifact someone asks for every Monday is a candidate for a routine. A routine that should start on its own and prepare an action for someone else is a candidate for an agent, and only then is the extra governance worth it.

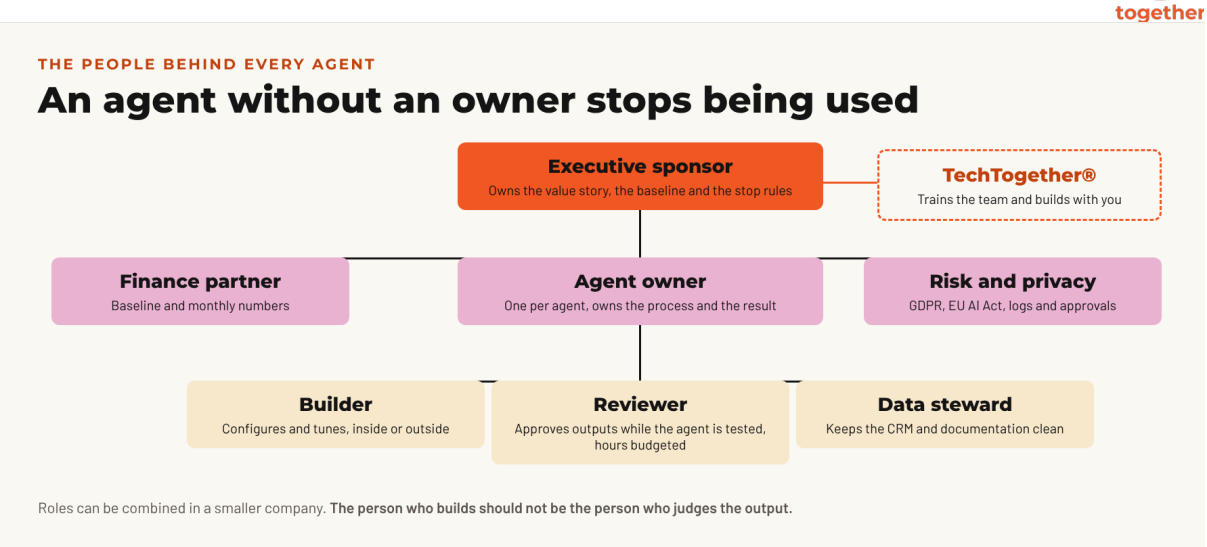


The people behind every agent

An agent is a new piece of operating capability, and like any capability it needs people around it. In a large organisation these are separate roles. In a smaller one a person may hold two of them.

- **Executive sponsor:** owns the value story, the baseline and the stop rules, and represents the programme at the board.
- **Agent owner:** usually the person who runs the workflow today. Writes the job description, decides what good output looks like and is accountable for the result. An agent without a named owner stops being used within weeks.
- **Builder:** configures, tests and tunes the agent. This can be someone inside who went through the training, an IT colleague or an outside partner working alongside them.
- **Reviewer:** approves outputs while the agent is being tested. This is the most underestimated cost in any agent programme, so the hours are budgeted.
- **Data steward:** keeps the sources clean, from CRM fields to the documentation the agent answers from.
- **Finance partner:** records the baseline and produces the monthly numbers, so the report to the board is prepared by finance and not by the people who built the agent.
- **Risk and privacy lead:** checks the approval steps and the logs against the GDPR and the EU AI Act.

One rule holds at any size: the person who builds the agent should not be the person who judges its output.



Three limits worth naming

A vision without limits ends in disappointment, and three limits are worth naming at board level.

- **The jagged frontier.** AI is excellent at some tasks and quietly wrong at others that look just as easy. That is why every agent keeps a human approval step.
- **Data.** An agent is only as good as the CRM, documentation or partner data it reads.
- **Autonomy.** Agents prepare and propose. People decide and send.

Within those limits there is a lot of room, and a board that names them up front earns the trust to approve a bolder plan.

The rules that apply

The EU AI Act was amended by the Digital Omnibus on AI, Regulation (EU) 2026/1744, which was published in the Official Journal on 24 July 2026 and entered into force on 27 July 2026. The main dates for an organisation that uses agents:

- **Since 2 February 2025:** no AI for prohibited practices such as harmful manipulation, social scoring or emotion recognition at work (Article 5). Organisations take measures to support the AI literacy of the people who work with AI (Article 4). The Omnibus softened the wording of this duty from ensuring a sufficient level to supporting its development, but the duty remains.
- **Since 2 August 2026:** people are told when they are dealing with AI, unless that is obvious, and AI generated images, audio and video that could pass for real are labelled (Article 50).
- **From 2 December 2027:** high risk uses listed in Annex III, such as recruitment, decisions about employees, consumer credit and the pricing of life and health insurance, need human oversight by competent people, logs kept for at least six months and workers

informed in advance (Article 26). Credit and insurance uses also need a fundamental rights impact assessment (Article 27).

- **From 2 August 2028:** high risk AI built into products covered by EU product legislation, such as machinery and medical devices (Annex I).

The GDPR applies today to every agent that touches personal data. In short: a lawful basis for each agent, only the data it needs, an updated privacy notice, a person who can change any decision with a significant effect on someone, processing agreements with AI vendors, lawful transfers outside Europe, access and erasure requests that also reach agent logs and memory, and breach reporting within 72 hours where required.

Dates checked on 7 November 2026. This is a starting point for the conversation with your privacy officer and legal counsel, not legal advice.

Questions for your next board meeting

- Has the foundation been installed for the people who use AI every day, whichever model we use?
- Has this board gone through it itself?
- Which moments in our business decide most of our revenue or cost, and who spends time on them today?
- Which agents score highest, and which line of the P&L will each one move?
- Which baseline is recorded before the build starts, and who owns it?
- What are the written thresholds and stop rules at each gate?
- Who reviews the output while an agent is tested, and how many hours does that take from their week?
- Where are the instructions, configurations and logs stored, and are they company property if we change supplier?
- Which board tasks could move to artifacts this quarter, so the board learns by using AI itself?

One to listen to

If you want one source that keeps the board conversation current without reading every report, try Azeem Azhar's Exponential View podcast. The episode *Why AI isn't showing up on your bottom line*, from 4 June 2026, speaks directly to the gap this paper starts from. It makes good company on the way to a board meeting.

What sets TechTogether® apart

TechTogether® is AI first. We are not an advisory firm that added AI to its services when the market started asking for it. We started with AI, and building agents is what we do every day, for our own work and for our clients. That shows in how we work:

- We build agents ourselves. Lavanya Trehan leads a team of twelve specialists at Crease Design Works who run the hands on builds.
- Both founders are in every session, so the people who designed the method are the people who answer your questions.
- Everyone works on their own laptop, with their own voice and their own work. Nobody sits and watches someone else click.
- The foundation we teach works whichever model your organisation uses, so you are not tied to one supplier.
- What we build stays with you. Instructions, configurations and knowledge files are company property, on your own laptops or in your own cloud.

A board that went first

Bizzmine, a Belgian quality management software company serving regulated industries such as life sciences and medical technology, was the first organisation to go through our training. In September 2026 we trained twenty people from their sales and marketing team and their C level board.

> *"We would highly recommend an inspiration session with Wendy and Lavanya to any organisation looking to take its use of AI to the next level. Their practical and engaging approach helps you see new possibilities, understand the value of AI agents and gain the confidence to start building them yourself."*
>
> **Alexander Werre, Chief Sales Officer, Bizzmine**

> *"The session with Wendy Morée and Lavanya gave our Marketing and Sales teams a practical understanding of how AI agents can support their daily work. We already use AI extensively at Bizzmine, but learning how to build our own agents was an important next step. The hands on format made the training especially valuable and gave both teams the confidence to start applying AI agents in their own workflows."*
>
> **Lara van Rijswijk, Chief Marketing Officer, Bizzmine**

Good training does not have to cost a fortune

We believe a good AI training should be within reach of every organisation, not only the ones with a consultancy budget. These are our rates for 2026, excluding VAT:

- **Private executive training:** €1,500 for four sessions of two hours over four weeks, with work between sessions. You choose which agent to build from our list, and afterwards you have 30 days of direct access to our engineers. Extra two hour sessions cost €375.
- **Team training:** €4,000 for one full day on site for up to fifteen people, covering LLM settings and privacy, persona and brand book, and a live agent build, followed by 30 days of support.

For a supervisory board and an executive board together, the team training day is the format we recommend.

About the authors

Wendy Morée is the founder and CEO of Moreequality and founder of TechTogether.AI. She is an executive and sales leader with twenty years in enterprise technology, including her time as Regional Vice President for Northern Europe at Salesforce, and sits on advisory and supervisory boards herself, including the advisory board of Frankly Solutions. She is an ambassador of NetworkC CIO Top Netherlands, was ranked #16 by Favikon in 2026 among the most influential female CEOs and tech leaders in the Netherlands, and was featured on NBC Washington as one of the Top 35 Global Women Thought Leaders (2024).

Lavanya Trehan is the founder of Crease Design Works in New Delhi, a design strategy studio working across branding, UI/UX and content, and delivery partner of TechTogether.AI. She is a TEDx speaker, trained at Central Saint Martins and the London School of Economics, and leads the hands on builds.

If you would like to talk any of this through, we are glad to.

Contact

Wendy Morée: wendy@techtogether.ai, +31 6 30 55 76 99

Lavanya Trehan: lavanya@techtogether.ai, +91 81301 34469

Website: www.techtogether.ai

Sources

- McKinsey and QuantumBlack, The state of AI in 2025: Agents, innovation, and transformation, November 2025.
- Bain and Company, Field notes from the generative AI insurgency, Global Private Equity Report 2025.
- Gartner, Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027, press release, 25 June 2025.
- Salesforce, State of Sales, 7th edition, 2026.
- Dell'Acqua and colleagues, Harvard Business School and BCG, field experiment on the jagged technological frontier, Organization Science, 2026.
- European Union, Artificial Intelligence Act, Regulation (EU) 2024/1689, as amended by Regulation (EU) 2026/1744.
- European Union, General Data Protection Regulation, Regulation (EU) 2016/679.
- Azeem Azhar, Exponential View podcast, Why AI isn't showing up on your bottom line, 4 June 2026.

© 2026 TechTogether®, a registered trademark owned by Moreequality. Written by Wendy Morée and Lavanya Trehan.