

— SURVIVAL GUIDE • 10 CHAPTERS

A Little *AI* Survival Guide.

For HR people, and everyone else. What is unfolding right now, and what you need to be aware of. No magic promises. Facts, laid end to end, and an invitation to verify them.

FREE TO READ AND DOWNLOAD • NO TRACKING, NO SIGN-UP.

WRITTEN BY GUILLAUME ALEXANDRE

/ LE BRASSUS • VAUD

/ [GATES-SOLUTIONS.COM](https://gates-solutions.com)

/ [LINKEDIN](#)

• BEFORE YOU OPEN


A guide, not a brochure.

This document gathers the ten chapters of a series built for one purpose only: looking squarely at what artificial intelligence is doing to work, and to your professions. No magic promises, no forty-nine-euro course. Facts, laid end to end, and an invitation to verify them.

Every number, every quote, every announcement points to its source, listed at the end of each chapter. The information is public. The only question that matters: once the picture is assembled, does it hold up?

The articles remain freely available. This treatise gathers the complete version — to read, keep, and share.

Read this before everything else



You've already come across three posts this week explaining, rocket emoji in hand, that AI is going to change everything, with a link at the bottom to a forty-nine-euro course. This text is not that.

This is a survival guide. The word is chosen carefully. Survival, because part of what you do today, in your teams, your processes, your job descriptions, won't make it through the year as it is. I'm not saying that to scare you. I'm saying it so we can look squarely at what is already happening.

Why me, and why now

In the past few months, I have multiplied my capacity to produce by twenty. Alone. I rebuilt several sites from the ground up, reconstructed the architecture of my business, set up an ERP, built tools that didn't exist. Even with a full team, we wouldn't have gotten there. Not better — not at all.

I've seen the power up close, hands in the engine. And I made a choice that was anything but theoretical. I was handed an opening to join a team chasing superintelligence. I said no, because I don't believe in what it's chasing. I could have helped speed up the machine. I chose the opposite: helping people keep their footing, and keep their hands on the wheel. That's where this series comes from.

The rules of the game

Everything I put forward here will be sourced. Every number, every quote, every announcement will point to its source, and you'll be able to check for yourself. The information is available, public, one click away. I'm not asking you to believe me. I'm asking you to look at the same facts I'm looking at.

And don't ask me whether I'm qualified to talk about this. That's not the question. I'm not showing up with any title to wave around. I lay out facts, and I put end to end information anyone can consult. The only thing that matters is whether the picture, once assembled, holds up. Read it, and judge that.

It's open to debate, of course. But not just any debate.

I want to avoid the trap that sterilized environmentalism. That false symmetry where, on one side, you feel guilty about the water running while you brush your teeth, and on the other, the planet keeps getting torched at industrial scale. That kind of debate only serves to clear consciences. It keeps people busy, it reassures them, it changes nothing. Here, we're not going to tear each other apart over whether you should use an AI to write your emails. We're going to look at what is really at stake, and at what scale.

The point to remember before all the others

What's coming is neither a hypothesis nor an opinion. The plan is written. It is even published. And it is being rolled out, whether you want it or not.

When an AI lab raises four billion to set up a deployment company with consulting firms as shareholders, that's not a rumor, it's a press release. When another one buys a company to acquire, in one move, a hundred and fifty engineers whose job is to go set up inside client companies, it's written in black and white in the job description.

The plan isn't hiding. That's the dizzying part. It's right in front of you, and almost no one reads it.

Three examples of what we'll look at together

Six weeks. That's the time between two major versions of the same large AI model, this spring. Not six years.

Six weeks. That's what an exponential looks like up close.

A trillion euros a year. That's the shortfall the head of one of Europe's rare labs came to describe before the French National Assembly, the equivalent of ten percent of the continent's payroll. The most worrying part isn't the number, it's the level of the questions he was asked.

Four billion dollars. That's the stake put down to create a company that no longer sells software, but comes straight inside your walls to do the work, and stays until it runs.

All of it is sourced. All of it will be detailed.

The table of contents

Here are the ten steps of this guide. We'll move in this order, because it tells a story that holds together.

1. Read this before everything else. You are here.
2. The curve won't wait for you. The exponential, and why it hurts.
3. You are not hidden behind your ERP. What MCPs are, and why they concern you.

4. Salary plus tokens equals an engineer ten times more productive. The new value equation.
5. Developers are the first copyist monks. The printing press is coming.
6. And the next scriptorium is yours. Why HR and everyone else come next.
7. The Forward Deployed Engineer enters the building. The job that sets up shop inside your company.
8. They're not after your software, they're after your payroll. The mechanics of the alliances.
9. Who captures the value, and who picks up the tab. The wealth shift, and its social blind spot.
10. Survival runs through local. Taking back control of your tools and your data.

The last point is also mine. What I do with all of this has a name, and we'll get to it without turning it into a pitch. For now, only one thing matters: becoming aware.

You can keep thinking you have time. Or you can read on.

I'm a sourcer. If you look for me, you'll find me.

The curve won't wait for you



People throw the word exponential around until it stops meaning anything. So let's put down a number, a verifiable one, and look it in the face.

The curve has a measurement

An independent research institute, METR, had a simple idea. Instead of grading AI with exam scores, measure the length of the tasks an AI can carry through to the end on its own. The yardstick: the time it takes a human professional to accomplish the same thing.

The result fits in one sentence. That length doubles roughly every seven months, and it has done so steadily for six years. At the end of 2022, when ChatGPT came out, the best models could hold a thirty-second task. By early 2026, some complete, on their own, tasks that take an expert several hours. And the trend isn't slowing, it's accelerating: over the recent period, the doubling happens more like every four months. The original measurement dates from spring 2025, and every generation released since has only confirmed the slope, or steepened it.

×2

The length of autonomous tasks doubles every ~7 months, and every 4 months recently.

30 s → h

From thirty seconds in late 2022 to several hours of expert work in early 2026.

×20

My capacity to produce, alone, multiplied by twenty in a few months. Not the tool — the overhaul of how I work.

That's what an exponential is. Not a slope you climb calmly at your own pace. A step that moves while you're deciding whether to take it. By the time your steering committee delivers its opinion on AI, the capability it was debating has doubled. By the time you draft the usage charter, it has doubled again.

What it looks like from the inside

I'm not telling you about a report I read. I'm telling you about what I lived. In the past few months, I multiplied my capacity to produce by twenty. Alone. Several sites rebuilt from scratch, a business architecture reconstructed, an ERP installed, tools created from nothing. Even with a full team, we wouldn't have kept up. Not better — not at all.

Want a concrete benchmark for the speed? A focus group I ran four months ago for a research project on HR already feels like it belongs to another century, given how much the tools have changed since. And this guide you're reading, from the first brainstorm to going live, layout included, took me less than an afternoon. Let's be honest about what that means: I'm talking about writing and laying out a body of thinking, not conjuring it out of thin air. A large share of the sources comes from groundwork that's been running on my side for three years. AI didn't think in my place. It collapsed the time between the thought and its finished form. That's the curve, at the scale of a single person.

Now, the honesty I owe you

And it runs against what you'd expect from a text about power.

The same institute, METR, ran another study. A randomized controlled trial, the protocol usually reserved for drugs. Seasoned developers, on their own code, projects they knew by heart. Verdict: with the AI tools of early 2025, they took nineteen percent longer. Slower. And that's not even the most troubling part. Those same developers were convinced they had been twenty percent faster. Nearly forty points of gap between what they felt and what had actually happened.

One caveat is in order, and it's a delicious one. That study was run with early-2025 tools. On AI timescales, that's an eternity. And I'm not the one saying so, METR is. The same institute re-ran its protocol, and in early 2026, on the same developers, the sign flipped: instead of nineteen percent of time lost, they now estimate a gain of around eighteen percent. The reason is plain: agentic tool usage exploded in 2025, and a tipping point came in late 2025 with the arrival of Gemini 3, then Claude Opus 4.7 and GPT-5.5, which made a spectacular leap on code. To the point where a growing share of those developers would now refuse to work without them. There's the curve, caught red-handed: a landmark study overturned by its own author within a year.

Let's stay honest, because that's what separates a fact from a sales pitch. METR says it themselves: these new numbers are fragile, selection bias, small sample, and they're reworking their method. And the recent studies, taken together, sketch a wide spread, from outright slowdown to massive acceleration depending on the developers, the language, and how clean the code is. To my mind, big studies like this one are already obsolete by the time you read them, because everything has changed in six months. So don't hold on to the percentage. Hold on to what doesn't expire: the gap between what you feel and what is actually happening. No new model version will ever erase that, because it isn't a matter of tools, it's a matter of humans.

So, does AI make you slower, or does it multiply you by twenty? Both. And that is the whole subject of this guide.

A French creator, Micode, framed the risk with a line that stuck: AI won't replace us, it will make us stupid. His thesis, in the video La Fabrique à Idiots, isn't that the machine takes your job. It's that it takes your effort. Delegate without thinking and you atrophy, until you end up mistaking the text produced for the skill acquired. That is exactly the trap the perception gap reveals: you feel stronger while you're getting weaker.

The lever isn't in the tool

The tool alone does nothing. Bolted onto an old way of working, it slows you down, because it adds mental load and constant back-and-forth. You spend your time explaining, rereading, correcting. That is exactly what the study measured.

My x20 didn't come from a plugin I switched on one morning. It came from a complete overhaul of how I produce. Of what I delegate and what I keep. Of how I chain tasks, how I break down a problem, how I check a result.

*The lever isn't in the tool. It's in **the hand that holds it**, and in the head that decides what to entrust to it.*

That is precisely what separates the one falling behind from the one multiplying by twenty. The same tool. Two worlds.

The whole difference is there, between a crutch and a lever. The crutch carries you, spares you the effort, and ends up atrophying you. The lever always demands your strength, but it multiplies it. Micode describes the crutch. I'm talking about the lever. It's the same object, and they are two fates.

This chapter isn't a threat, it's a filter


The curve will keep going, whether you're ready or not. No one will pause it while you get organized. The only thing still yours is choosing which side of the lever you'll be on.

And that, no tool will ever decide for you.

● SOURCES

-
- 01 **MARCH 2025** · METR, Measuring AI Ability to Complete Long Tasks, March 2025 (doubling every seven months over six years): metr.org
-
- 02 **2025** · AI Digest, A new Moore's Law for AI agents (recent acceleration toward four months): theaidigest.org
-
- 03 **JULY 2025** · METR, Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity, July 2025 (nineteen percent slower, perception gap). A study deliberately presented as dated: early-2025 tools, outpaced since the late-2025 shift, which serves the demonstration of the curve. The lesson to keep is the perception gap, not the number. metr.org and arxiv.org
-
- 04 **FEBRUARY 2026** · METR, We are Changing our Developer Productivity Experiment Design, February 2026 (on the same developers, the estimate goes from -19% time lost to a gain of around 18%, explosion of agentic tools in 2025, caution about selection bias): metr.org
-
- 05 **2026** · Ingo Eichhorst, State of AI Coding Efficiency 2026, meta-analysis (spread of results from slowdown to acceleration depending on studies, languages, code cleanliness): ingoeichhorst.medium.com
-
- 06 **2025** · Micode, La Fabrique à Idiots, YouTube (AI as a cognitive crutch, the risk of atrophy): youtube.com

You're Not Hidden Behind Your ERP



There's a quiet belief, widespread among HR leadership, and it's going to cost you dearly. It fits in one sentence: my job is protected, because my data and my processes live inside a big locked-down piece of software — my ERP, my HRIS, my ATS. A fortress. Nobody touches it without the vendor.

That fortress has just lost its walls. And the culprit has a three-letter name: MCP.

What an MCP is, without the jargon

Think back to when every device had its own charger. One cable per brand, a drawer full of useless wires. Then USB-C came along, one single plug, and everything connected to everything.

MCP is the USB-C of artificial intelligence. A single standard that lets you connect a piece of software, through its technical interface, directly to an AI. In plain terms, you plug your system into a large language model, and the AI can now read, write and act inside that system, in everyday language, simply because you asked.

This is no lab promise. In eighteen months, MCP became the standard for the entire industry. Launched by Anthropic in late 2024, it was adopted by OpenAI, Google, Microsoft and AWS, then handed over to an independent foundation so that no one would own it. The development kit went from one hundred thousand downloads a month to ninety-seven million. When tech's fiercest rivals all adopt the same plug in under two years, that's not a fad. It's a tipping point.

100k → 97M

Monthly downloads of the MCP development kit, in eighteen months.

18 months

To become the de facto standard, adopted by all of tech's rivals.

1/2 day vs 20 years

Cleaning 20,000 contacts: half a day, where twenty years hadn't been enough.

What it changes, as told by my weekend

Let me give you an example I just lived through, one you'll be able to verify by the time you read these lines.

I work on Jarvi, my candidate management software. Jarvi just released its MCP. In half a day — a single one — I started cleaning a database of more than twenty thousand contacts. Putting phone numbers back into the right format. Overhauling the entire naming scheme. Bringing order to chaos accumulated over years.

Ask yourself, if you're in the business. Who in recruiting can brag about having a truly clean database? In twenty years, I never managed it. Not once. Now it's happening, and for the first time in my career I'm going to have a clean candidate database, GDPR-compliant, partly self-cleaning, and working. By the time this article comes out, it will be done.

Half a day versus twenty years. That's the universal plug plugged into the right trade.

The wall becomes a door

That's why your fortress has changed in nature. Yesterday, your software locked you inside its features. If the vendor hadn't planned a feature, it didn't exist, period. You waited for the next release, sometimes for years.

With an MCP, that wall becomes a door. Any process in your company, the moment you can state it clearly, you can automate directly inside the system. And even if your software never planned for this or that feature, you can add it on top. The tool's capability is no longer decided by the vendor alone. It's decided by whoever can articulate what they want.

There's good news in there, and a threat. The good news: you can finally make your systems do what you'd been dreaming of. The threat: someone else will do it in your place if you don't.

Where this becomes my job

And this is where I start pushing further. I now work directly with ATS vendors — the recruiting software makers — to see how far we can go. Improving metrics, building dashboards the software can't produce natively, connecting the data to external AIs. Because we now have access to the data, and by adding a recruiter's knowledge of the craft, we should pull off things an ATS, by its very nature, can't do.


I don't know yet whether it deserves to be called a skill. But one thing is certain: tackling problems that were strictly impossible yesterday, as long as the vendor hadn't shipped them, doesn't scare me. Quite the opposite.

The realm of the possible just blew wide open. Enormously. The only question left is who's going to rush in. Your software vendor. An outside contractor. Or you.

● SOURCES

-
- 01 2026 · TechAhead, Model Context Protocol, the Enterprise AI Integration Standard (the USB-C of AI, OpenAI adoption March 2025, Google April 2025, donation to the Linux Foundation December 2025, from one hundred thousand to ninety-seven million monthly downloads): techaheadcorp.com
-
- 02 2026 · WorkOS, Everything your team needs to know about MCP in 2026 (de facto standard, Linux Foundation governance, OpenAI and Block as co-founders): workos.com
-
- 03 2026 · Digital Applied, MCP Adoption Statistics 2026 (verified figures, official registry close to ten thousand servers, caution about inflated numbers): digitalapplied.com
-
- 04 2026 · Wikipedia, Model Context Protocol (OpenAI and Google DeepMind adoption timeline): en.wikipedia.org

Salary Plus Tokens



Ask yourself a simple question, you who recruit, who evaluate, who set compensation. What is an employee worth? Until now, the answer fit on one line: their salary. Tomorrow, it will fit in a sum: their salary, plus the tokens they consume.

That sentence sounds abstract. It has already become company policy.

The new payroll line already exists

In March 2026, Nvidia's boss, Jensen Huang, proposed giving his engineers an allowance of AI tokens worth nearly half their base salary. And he pitched it as a recruiting argument. Why be tempted by a signing bonus elsewhere, he said in essence, when here you get power.

Read carefully what that says. One of the most powerful executives on the planet now considers that an engineer's value is no longer just what they can do, but what they can get the machine to do. The salary pays for the head. The tokens pay for the leverage. And the company is willing to fund both, because it knows the combination is what produces.

The asymmetry that changes everything

Here's why this equation is so brutal. The two terms don't cost anything alike.

A mid-level developer costs between seven and thirteen thousand francs a month, salary and payroll charges included. Running AI agents at full throttle to multiply that developer's output costs, depending on usage, from a few hundred to one or two thousand francs a month. The models themselves are dropping in price by fifty to seventy percent a year. In other words, for a fraction of a salary, you multiply what that salary produces. And the gap widens every year in favor of the tokens.

Some have understood this so well they've turned it into a competition. In Silicon Valley they call it tokenmaxxing, and star engineers burn through the equivalent of more than eighty thousand francs in tokens a month. One of them puts it bluntly: I probably spend more on AI than my own salary. In his eyes, that's not waste. It's an investment in his own leverage.

$\approx \frac{1}{2}$ salary

The token allowance proposed by Nvidia, pitched as a recruiting argument.

7-13k CHF

The monthly cost of a mid-level developer, where agents cost a few hundred francs.

140
CHF/month

My actual AI bill. Set against a single salary, the asymmetry hits you in the face.

What it looks like at my place, without a Silicon Valley budget

I don't have a hundred thousand dollars of tokens a month. Far from it. And I'm going to be completely transparent, because that's the whole spirit of this guide. Here is my actual monthly bill. A Gemini Pro account attached to my Google Workspace, enough to run AI Studio and NotebookLM, eighteen francs. A ChatGPT Plus I've had since the beginning and barely use right now, mostly for images, twenty francs. And a Claude Max account, the five-times plan, one hundred francs, with which I do ninety percent of my projects. Total, about one hundred and forty francs a month.

One hundred and forty francs. Put that figure next to a single salary, and the asymmetry hits you in the face. This isn't a thought experiment, it's my bank statement.

That's the equation, seen from a small operation. My time, plus one hundred and forty francs of power, against the payroll of an entire team. And the result tips the way nobody expected.

And it doesn't stop at developers

People often assume this story only concerns technical jobs. Let me tell you another one.

A few months ago, to prepare a talk to HR people about what AI can really do, I built a complete marketing strategy for a new market segment. Product visuals, a campaign draft, a slogan, social media ads. All of it in forty-five minutes, with AIs, on the wifi of a TGV doing three hundred kilometers an hour.

I'm not a marketer. And I'm not claiming for a second that a team of marketers wouldn't have done better than me. They would have, no question. But here's the real question, the one that matters to a company. While a team without AI polishes one concept in a week, I, alone with my tokens, could have tested fifty. The quality of one finished concept versus the speed of exploring fifty leads. For leadership, that trade-off is settled fast.

Extrapolate. One salary plus tokens moves faster, in the exploration phase, than five or even ten salaries without AI. Yesterday that was only true for developers. It's becoming true for marketing, communications, analysis, legal, and the rest.

The link that snaps, and nobody wants to name it

Of course, exploration is not production. In a real company, there are bottlenecks slowing all of this down. But look at where those bottlenecks often come from. They don't come from the technology. They come from the chain of command, the layers of sign-off, the management.

And what's happening right now, at this very moment? Those layers are being removed.

The phenomenon has a name, the Great Flattening. According to Gartner, one organization in five will use AI to flatten its structure by 2026, cutting more than half of its middle management positions. And the most brutal example is brand new. In May 2026, Meta notified eight thousand layoffs, about a tenth of its workforce, and canceled six thousand open positions — nearly fourteen thousand jobs hit in one week. The trigger: AI capable of executing and coordinating what those managers used to coordinate.

But the most chilling part isn't the number. The very day the notifications went out, a recording leaked. In it, Zuckerberg explains, during an internal meeting, that his AI models learn by watching really smart people do their jobs. Follow the sequence carefully: the company tracked the activity on its employees' computers, with no way to opt out, to train its AIs on how they work, then laid people off. Learn from the best, then let them go. In the same recording, he admits it was not in the company's strategic interest to explain frankly to the teams how the program worked. The employees understood it just fine: more than a thousand signed an internal petition against the program.

Take in the irony, because it's cruel. For decades, we promoted the best, the most senior, the most brilliant, into management positions. It was the royal road, the only recognized way to climb the hierarchy. We spent years pushing talent toward the middle of the pyramid. And it's precisely that middle that's now being removed.

Let's be honest about the downside, because it's heavy. Cutting out the middle is not painless. Thirty-seven percent of employees already say they feel lost without managers, and nearly half of executives doubt they can manage what's left. You're also removing the drive belt, the mentorship, the next generation. Flattening without thinking means creating a void where culture and judgment used to flow.

But the direction of the movement leaves no doubt. The pyramid is being squeezed through the middle.

The line HR will have to learn to read

If you work in human resources, here's what's landing on your desk, whether you want it or not. Soon, comparing two candidates for the same position will no longer be done at equal salary. It will be done at equal leverage. The real question will no longer be how much they cost, but how much they produce per euro of tokens spent.

And that's a reading grid no current job description contains. Nobody, in your competency frameworks, has a box for that.

The nuance, because people will lie to you about it

You'll be sold the equation as a magic formula. Buy tokens, get an engineer ten times more productive. That's false, and chapter 2 of this guide explained why. Tokens multiply no one on their own. Plugged into someone who hasn't changed the way they work, they burn money without producing anything. It's been measured: entire teams have watched their AI bill become the second-largest expense after salaries, with no measurable productivity gain.

The leverage isn't in the tokens. It's in the person who knows what to ask of them.


So the right equation isn't salary plus tokens equals ten times more. It's salary, plus tokens, plus the skill to use them, equals ten times more. And that skill is the only one of the three you can't buy by the million tokens.

It's also the only one still in your hands.

● SOURCES

-
- 01 **MARCH 2026** · Futurism, Bosses Are Blowing More Money on AI Agents Than It'd Cost Them to Just Pay Human Workers (Jensen Huang proposes tokens worth half of base salary as a recruiting tool, tokenmaxxing, more than one hundred thousand dollars a month, quote from the Stockholm engineer): futurism.com
-
- 02 **2026** · amux.io, AI Coding Agent Costs in 2026 (a mid-level developer costs eight to fifteen thousand dollars a month, lower agent costs, model costs dropping fifty to seventy percent a year): amux.io
-
- 03 **MAY 2026** · The Main Thread, AI Coding Break-Even (token prices verified as of May 21, 2026, implementation is only half the work, caution about any single productivity figure): the-main-thread.com
-
- 04 **2026** · LeanOps, AI Agents Burn 50x More Tokens Than Chats (the AI bill becomes the second-largest expense after salaries, gains not guaranteed without optimization): leanopstech.com
-
- 05 **MAY 2026** · Meta, May 2026, eight thousand layoffs, six thousand canceled open positions, leaked Zuckerberg audio on training AIs by watching employees, surveillance program with no opt-out, seven thousand reassigned to AI, internal petition. TechTimes (techtimes.com), eWeek (eweek.com), Common Dreams (commondreams.org).
-
- 06 **2026** · Fast Company, What happens to middle management when AI flattens your organization (the Great Flattening, up to 20% of firms cutting middle management by end of 2026, agentic AI as the trigger): fastcompany.com
-
- 07 **2026** · People Managing People, The Great Flattening (Gartner, 20% of organizations cutting more than half of middle management positions by 2026, the Amazon case): peoplemanagingpeople.com
-
- 08 **2026** · Lepaya, The Great Flattening (a 6.1% drop in the number of managers between 2022 and 2025, Meta Amazon Google Intel, 37% of employees feeling lost, mentorship and judgment as counterweights): lepaya.com

Developers, the First Copyist Monks



Before the printing press, copying a book was a trade. In the silence of monasteries, copyist monks spent their lives reproducing texts by hand, letter after letter, one page a day on a good day. It was a rare craft, respected, and well paid for its time. Then Gutenberg set up his press, around 1450, and within a generation the trade was gone.

But look closely at what happened. The copyist wasn't erased from the world. He moved. Those who understood printing became printers, publishers, proofreaders. Those who clung to their quills became relics. The knowledge didn't disappear, the gesture changed.

Today, the first scriptorium to receive its press is software development.

The code printing press is already here

Anthropic's boss announced it point-blank: AI will soon write the vast majority of code. And indeed, AI agents no longer settle for completing a line. They open a project, write an entire feature, test it, fix it, and ship it. The hand that held the quill now holds something else.

And it's already no longer a prediction. In May 2026, before the French National Assembly's commission of inquiry, Mistral's boss, Arthur Mensch, said something dizzying: at his company, engineers no longer write lines of code. The developer is no longer the craftsman who writes, he has become a manager directing agents. Remember the date, because it matters in a field where everything moves fast: this tipping point isn't announced for tomorrow, it's already behind us. We'll come back to that hearing, it says far more still, but for now let's hold onto that one sentence.

Everywhere, the same funereal conclusion is being drawn: the developer's trade is doomed. And everywhere, it's wrong, exactly as wrong as announcing the end of the written word in 1450.

0 lines

At Mistral, engineers no longer write lines of code (Arthur Mensch, May 2026).

150,000 lines

Of AI-generated code before one senior developer hit the wall.

9 projects · 6 weeks

What I ran in parallel, learning hands-on, not in a training course.

Developers are, for now, the big winners

It's the paradox few dare to state. Developers are the first ones hit, and at the same time the first winners. For a very simple economic reason: theirs is the only trade whose productivity gain translates immediately into money.

Senior coder Simon, in a recent video, says it with the same image I use here, Gutenberg's. His thesis matches mine: mastering AI in 2026 is the most powerful lever to accelerate a developer's career, not to doom it. The trade isn't dying, it's moving up a notch. From writing to supervising, from execution to intent. The copyist becomes a printer. It's exactly what Mensch describes when he says his engineers have become managers of agents: not the end of the trade, its shift upward.

But that climb is anything but automatic, and this is where I want to be honest with you.

The flip side, told by those who got burned

A senior developer publicly told the story of his disenchantment, after one hundred and fifty thousand lines of AI-generated code, with the tools from before the late-2025 shift, the one brought by Gemini 3, Claude Opus 4.7 and GPT-5.5. At first, the euphoria: he was shipping in a week what a team produced in a month. Then the wall. On a real production feature, the AI started looping, modifying seven files at random, stacking up bugs. His conclusion: the tool doesn't truly reason, and you have to reread every line before accepting it.

Since then, the models have crossed a threshold, and some of those limits have receded. But don't take away the wrong lesson. What remained true yesterday remains true today: the tool amplifies, it doesn't decide. The more powerful it gets, the more the question shifts toward whoever is steering it.

The two accounts don't contradict each other. They describe the two possible fates of the same trade. The one who learns to run the press prints thousands of pages. The one who believes the press thinks in his place produces thousands of wrong pages.

The press amplifies. It amplifies talent and error alike.

Why this chapter concerns you, even if you don't code

You may be telling yourself this is all a matter for IT people. That's a mistake, and it's the whole point of the chapters that follow.

If developers are hit first, it's not because they're the most fragile. It's because they're the best positioned: their work is the easiest to convert into machine output, and the easiest to measure in value. They are the special case that announces a general law. Because what's happening with code, the transformation of human work into something the machine produces for a few euros of compute, won't stop at code.

And above all, don't go looking for a training course

Here's the real question. How did the developers who are making it become good at this game? Not by signing up for a training course. By getting their hands dirty.

And this is where I want to warn you, HR people, because this wave will reach your professions, and I already know your first reflex. When you see things moving, you look for a training course. It's the reflex of an entire professional life. This time, it won't be enough.

Do the math. By the time a training body designs a program, gets it approved, and trainers themselves become good enough to teach it, the content is already outdated a thousand times over. The subject moves faster than any curriculum can be written. You'd be learning last year's state of the art, presented as a novelty.

I know because I'm living it. In six weeks, I've run nine projects in parallel. A site that scrapes data in real time from multiple sources to keep my community informed. The complete redesign of three sites — my business, my training program, a tool I co-run. Setting up an ERP to handle all the admin and sales for my sites. The full overhaul of an e-commerce site, in progress. A finished site for a friend's business. A complete ecosystem to manage meals and groceries at home. And this guide. Not to mention tools for recruiters and HR, also in the works. I'm not asking you to take my word for it, go browse my sites and judge the result.

And what I do on the latest one has nothing in common with the first, even though six weeks have gone by, not six years. I learned by doing: by understanding how a tool's capabilities work, by testing relentlessly, by pitting AIs against one another to see which understands best. If I had to freeze all of that into a training course today, it would be impossible. It would expire before it was finished.

What works is the opposite of a training course. Carve out time. Take on a real project head-on. Identify a concrete problem and try to solve it. In a safe sandbox, so nothing vital breaks, but try for real. Accept that it doesn't work now, dig into why, and watch it work five minutes later. It's all about iteration. You don't receive the recipe, you find it.


I can coach, mentor, push someone who's already elbow-deep in the grease, tell them have you seen this, have you tried that. But explaining from scratch everything that's doable to people demanding a final recipe without ever wanting to iterate — that, nobody can do. Because it doesn't exist.

That's exactly what we'll see further on. The next scriptorium to receive its press, after the developers, may well be yours. And the day it does, don't go looking for a training course. Get your hands dirty.

● SOURCES

- 01 MAY 2026 · Senior coder Simon, YouTube channel Codeur Senior, video of May 28, 2026, sponsored by TestSprite (the Gutenberg analogy, AI agents as code printing presses, developers as AI's big winners, mastering AI as a senior career lever): [youtube.com](https://www.youtube.com)
-
- 02 2025 · theSeniorDev, Why I Stopped Using AI as a Senior Developer After 150,000 Lines of AI-Generated Code (initial euphoria then limits, loops, production bugs, the need to reread every line): theseniordev.com
-
- 03 2025 · On the share of code written by AI and the trajectory toward nearly all of it: public statements by the heads of Microsoft, Google and Anthropic, widely reported. To be cross-checked against a recent primary source at publication time, as the figures move fast.
-
- 04 MAY 2026 · Arthur Mensch (Mistral), hearing of May 12, 2026 at the French National Assembly: engineers no longer write lines of code, the developer becomes a manager of agents, productivity gain doubled in six months, some jobs all but disappearing, value shifting from labor to capital. LCP (lcp.fr), Contrepoints (contrepoints.org).

And the next scriptorium is yours



The previous article ended on a promise, or a threat, depending on your mood. After the developers, the next workshop to get its printing press might be yours. The “might” was a courtesy. Let’s drop it.

This is no longer an IT matter

The numbers no longer talk about code. The International Monetary Fund estimates that about 40% of jobs worldwide are exposed to AI, and nearly 60% in advanced economies like ours. Goldman Sachs puts the share of work exposed to automation at the equivalent of three hundred million full-time jobs. And the World Economic Forum, in its benchmark report, projects that 41% of employers plan to cut headcount where AI automates tasks, and that 40% of the skills in demand will change within five years.

40-60%

Of jobs worldwide exposed to AI, nearly 60% in advanced economies (IMF).

+78M

Net job balance by 2030: 170 million created, 92 million destroyed (World Economic Forum).

281,000

Participants in the largest AI hackathon, at TCS. The right answer is not a training course.

Exposed does not mean eliminated, and I'll come back to that in a moment, because that's where the lying is heaviest. But note the direction. The wave doesn't stop at the IT department. It's rising toward marketing, finance, legal, analytics, and yes, human resources.

And since this is my trade, let me add a practitioner's note, because it applies to every one of your functions. Not every use of AI is equal, and the dividing line fits in one word: data. Who owns it, and who does it serve. Take recruiting. You're being sold automated sourcing built on data trapped inside LinkedIn. But that data isn't yours, LinkedIn's internal AI pursues its own goals, not yours, and the supposedly magical solutions that claim to access it only survive thanks to data harvested through means whose compliance could raise questions. The whole point of my training program, peopleattractiontheory.com, is to dismantle that mirage. Conversely, when I take the data I truly own, structured in my own recruiting software, Jarvi, and I can now query it with precision thanks to the MCPs we saw in chapter three, then it's a tremendous lever. The same difference as between a crutch and a lever: borrowing someone else's data makes you dependent, mastering your own makes you powerful.

The optimistic objection, and why it misses the point

At this stage, a voice always rises, and it's a respectable one. The voice of those who remind us that AI will also create jobs, that new professions will be born, as with every technical revolution.

The most illustrious person saying this is Yann LeCun, one of the fathers of modern AI. Let me be blunt: I admire him. Scientifically, without him and the lab he built in Paris, there would probably be less chance of a European champion like Mistral existing, since two of its three founders came out of it. When he says tomorrow's work will be about augmentation, that everyone will become the conductor of their own team of AIs, and that entirely new professions will appear, I believe he's right.

I even have a confession to make. I was lucky enough to spend some time close enough to that world to find myself taking back-to-back calls in the room next to his. I crossed paths with him. And I, who am not exactly known for holding my tongue, didn't dare speak to him. Pure fanboy effect. I mention it because it shows where I'm speaking from: not from a pedestal, but as a guy sincerely in awe of what he has accomplished.

And that's precisely why I can contradict him without the slightest animosity. On the destination, I'm with him. Not on the path.

The real problem is the tempo

The World Economic Forum report actually proves him right on the raw numbers. One hundred seventy million jobs created by 2030, against ninety-two million destroyed. A net positive balance of seventy-eight million. On paper, everything is fine.

Except that a net balance masks a brutal asymmetry of speed. Deciding not to replace a departure, because three well-equipped people are now enough, takes one meeting. Deciding to cut positions takes one meeting. But creating professions at scale, defining them, writing the job descriptions, training people, hiring them, getting them settled in — that takes years.

The cut is immediate. The creation is slow. Between the two, a gap opens up, and that gap is where people fall.

The balance may well be positive in 2030. It says nothing about what the people for whom the cuts arrive years ahead of the creation will go through between now and then.

The example is right before our eyes, and it's ironic. While we're being promised new professions, Meta cut thousands of jobs to fund its AI infrastructure — we saw how in the previous chapter. The new professions, we're still waiting for. The cut jobs, they're already gone. That's the tempo. And the tempo is the real subject for human resources. Not the 2030 horizon, but the three years leading up to it.

Besides, the executives themselves don't agree. Anthropic's CEO warns that AI could wipe out half of entry-level white-collar jobs within one to four years. Nvidia's CEO argues, on the contrary, that it will create more highly skilled jobs. When the people building the technology diverge that much, beware. Nobody holds the truth, and that's exactly why you must act without waiting for them to settle it.

The honesty I promised

Exposed does not mean doomed. The figures measure automatable tasks, not actual layoffs, which are far lower. A task is not a job, and a job is not a person.

But don't reassure yourself too quickly. Between the intact job and the eliminated job, there's a third category the statistics don't count: the job emptied of its substance. Kept, but transformed. The one from which the meaningful tasks have been removed, one by one, until only a shell remains. That's often where it starts, well before the layoff.

So what do we do, without panicking

Faced with this, there are three wrong answers and one right one.

The first wrong answer is to close your eyes. We've covered that — the curve won't wait for you.

The second is to run off and find a training course. We've covered that too. By the time it's written and approved, it's already obsolete. You don't learn this in a classroom, you learn it with your hands in the machine.

The third is to cut fast to save money. It's the most tempting, and the most dangerous. In the final chapter we'll see how a company that was nonetheless fanatical about AI got burned doing it, then had to rehire.

The right answer fits in an old word brought back into fashion. The hackathon. But not the developers' kind. A hackathon for the company's other functions. The principle is simple and powerful. You take a real problem you don't know how to solve, you give yourselves a limited time, and you search. Together. No process, no infamous "that's how we've always done it."

This isn't a whim of mine. Deloitte observes that the vast majority of companies remain stuck at the experimentation stage, unable to move their AI ideas into production, and presents the internal hackathon as one of the fastest ways to close that gap. Canva shuts its offices for an entire week to do this, and gets features that actually ship in its products. The IT services firm TCS gathered more than two hundred eighty thousand employees from fifty-eight countries in a single AI hackathon. And, closer to my point, a Belgian government agency organized a hackathon where teams worked locally, on secured machines, with open models and their own data, to demonstrate that sovereign AI can improve internal processes without ever letting information leak.

That's what needs to be done. Set the plasticity of your teams' brains back in motion. Get them out of the corridor they've been walking down for ten years. Give them a sandbox, a real problem, and the right to fail several times before succeeding. That's how you learn now. No other way.


One question remains

To lead all this, you need someone who knows. Someone who walks into the company, observes the processes, and shows where AI can fit in without breaking everything. That someone exists, they carry a title still little known, and they're the subject of the next chapter.

● SOURCES

-
- 01 2025 · IMF, about 40% of jobs worldwide exposed, nearly 60% in advanced economies, and the Amodi versus Huang debate, via World Economic Forum ([weforum.org](https://www.weforum.org)).
-
- 02 2025 · Goldman Sachs, equivalent of 300 million full-time jobs exposed, and the exposure-versus-actual-layoffs nuance, via SQ Magazine (sqmagazine.co.uk).
-
- 03 JANUARY 2025 · World Economic Forum, Future of Jobs Report 2025 (170 million created, 92 million destroyed, net positive balance of 78 million, 41% of employers planning to reduce headcount, 40% of skills changing): [weforum.org](https://www.weforum.org)
-
- 04 2025 · Yann LeCun, Génération Do It Yourself interview, episode 543 (augmentation, new professions, everyone the conductor of their own AI team): gdiy.fr
-
- 05 SEPTEMBER 2025 · The FAIR-to-Mistral link, two of the three founders from Meta, the third from DeepMind. TechCrunch (techcrunch.com), IBM (ibm.com).
-
- 06 2026 · Meta cuts eight thousand jobs to pay for its AI infrastructure. Tom's Hardware (tomshardware.com).
-
- 07 2025 · Internal AI hackathons, Deloitte on stalled pilots and the hackathon as an accelerator, open to non-technical staff. AngelHack (angelhack.com), Devpost (info.devpost.com).
-
- 08 MAY 2026 · Canva closes for a week of AI upskilling and its hackathon. HR Grapevine (hrgrapevine.com).
-
- 09 2026 · TCS, more than 281,000 participants in the world's largest AI hackathon. TCS (tcs.com).
-
- 10 MARCH 2026 · Local, sovereign AI hackathon, AI Hackathon 2026 organized by the Belgian Digital Transformation Office (BOSA), in Brussels March 17 to 19, 2026 as part of EU AI Week. Work done locally on secured machines, open-source models only, teams' internal data, RAG system, within Belgium's stated digital sovereignty strategy. European Commission page, Digital Skills and Jobs (digital-skills-jobs.europa.eu) and official BOSA page (bosa.belgium.be).

The Forward Deployed Engineer walks in



The previous article showed you the tool, MCP, that universal socket that opens up your systems. One question remained. Who holds the plug? Who walks into the company to decide where to plug in AI, and how?

That someone has a title, still unknown to the general public two years ago, and now becoming the most coveted job in tech. It's called the Forward Deployed Engineer.

The job that exploded in silence

The term comes from Palantir, which coined it more than ten years ago. The idea: take the engineer out of the office and send them inside the client company, in contact with the real processes, to build and deploy the solution on site. Not a consultant who hands in a report. Someone who moves in and delivers.

For a decade, nobody outside Palantir paid attention. Then AI arrived, and the job exploded. In one year, postings for this role jumped more than eleven hundred percent. OpenAI, Anthropic, Google, Databricks — all are building dedicated teams. Box’s CEO sums up the situation in one sentence: it will soon be one of the most sought-after jobs in tech.

+1100%

The explosion in postings for this role, in one year.

1 day/week

The time some teams can claw back from valueless, automatable tasks.

Palantir

The model invented there, now copied by OpenAI, Anthropic, Google, Databricks.

I owe you a clarification, because I don’t like operating under a mask. On my personal scale of values, few companies strike me as worse than Palantir. That’s an opinion, mine, and I own it. But it would be dishonest to let that aversion cloud my analysis. The model they invented, the engineer deployed at the client’s, will in my view become the norm. You can hate where an idea comes from and recognize that it’s going to win. That’s even the precondition for getting ready for it.

My conviction, and it is not the official definition

Let me be honest with you on one point, because it matters. The canonical definition of this job, inherited from Palantir, insists on code. A software engineer, first and foremost.

I believe AI is moving that cursor, and this is where I go out on a limb. The rising profile is no longer just the computer scientist. It's the person, sometimes from another field entirely, who became extremely competent in AI by teaching themselves, and who above all knows the company's processes from the inside. Their value isn't in the engineering degree. It's in the ability to look at a trade, spot what can be automated, and articulate it clearly enough for the AI to take it on. The trade outranks the code.

That's my thesis, not engraved truth. But I'm not pulling it from a book. I lived it. Remember my March assignment: it was a sourcer, an HR guy, they went and got to deliver an AI roadmap for the entire company. Not an engineer. Proof, by my own example, that what counts isn't the title on the business card, but the ability to understand a trade and retool it.

What I did in March, and what people asked me for

In early March, I carried out exactly this type of assignment, in a large SME in the Fribourg region. And note the starting point, because it says everything: I wasn't called in for recruiting or for HR, my original trade. I was brought in to walk into the company, talk with people, observe how things actually get done, and produce a roadmap on where to put AI into the processes. With considerable gains at stake. For some teams, we're talking about recovering the equivalent of one day per week. An entire day, spent on tasks with no added value, perfectly automatable.

And here's the detail I want you to remember, because it changes everything. I wasn't the one who showed up to hunt people down and replace them. It was the other way around. They were the ones who asked me. Here are the tasks that serve no purpose, they told me, and if I could get rid of them, I could finally focus on what moves the company forward.

Nobody is afraid of being relieved of what bores them. People are afraid of being replaced.

It's not the same thing, and that's the whole challenge of how these subjects are brought up.

I loved doing it. Sincerely.

The real work starts afterward

The report is the easy part. The hard part, and the decisive one, is going from the report to actual implementation inside the company. And that's where the shoe pinches, because companies don't have these skills in-house. They know they have to move, they don't know how.

That's why I'm making this bet: the Forward Deployed Engineer role is going to spread. First in large corporations, which have the means to go get one. Then, and this is where the real opportunity lies, in SMEs, which need it just as much and don't yet have access to it.

There's an insane opportunity here. For those who'll want to embody it, and for the companies that will know how to seize it. The question remains, once again, who will grab it.

- SOURCES

-
- 01 2026 · ZTABS, What Is a Forward Deployed Engineer (postings up more than eleven hundred percent in one year, quote from Aaron Levie, Box's CEO, dedicated teams at OpenAI, Anthropic, Palantir, Databricks, Scale AI): ztabs.co
-
- 02 MAY 2026 · MarkTechPost, What is a Forward Deployed Engineer, the AI Role OpenAI, Anthropic, and Google Are Hiring in 2026 (hybrid profile, technical plus business, applied AI skills): marktechpost.com
-
- 03 2026 · Wikipedia, Forward Deployed Engineer (Palantir origin, embedded at the client's, link with OpenAI's Deployment Company): en.wikipedia.org
-
- 04 2026 · Kore1, What Is a Forward Deployed Engineer (role less than two years old outside Palantir, profiles not strictly from engineering backgrounds possible): kore1.com

They're after your payroll



When a company senses that AI is getting beyond it, it has an old, reassuring reflex. It calls a big consulting firm. They'll tell us what to do, we'll sign, they'll install, and we'll be able to sleep. In 2026, that reflex is a trap. And to understand why, you have to look at what these firms have become.

Consulting is no longer neutral, it's an interested party

For decades, the big firms sold one thing: intelligence. Analysis, strategy, recommendations. That was their bread and butter, and they charged top dollar for it.

But remember the formula from Mistral's CEO, encountered earlier in this guide. Intelligence, now, is electricity converted into tokens. It's produced in near-unlimited quantities, for a fraction of yesterday's price. In other words, the product the firms sold at the highest price is becoming a commodity. They are threatened at the very heart of their business.

So they do what any rational, worried player would do. They keep selling you intelligence, but now their partners'. The kind that lets them stay in the loop and keep billing. The problem is that this choice is no longer neutral. And those aren't my words — the financial arrangements themselves say it.

The facts, and they are public

In May 2026, OpenAI launched a company dedicated to enterprise deployment, endowed with four billion dollars, financed by some twenty investors. Among them, consulting firms themselves: Bain, McKinsey, Capgemini. Read that carefully. The firms are no longer merely service providers for the tool they recommend to you. They are shareholders in it.

Right on its heels, OpenAI bought a company to pick up roughly a hundred and fifty deployment engineers in one stroke, those profiles who come settle inside your walls. The goal is explicit: capture the implementation margin on top of the model margin. Translated for you: the vendor no longer wants to sell just the engine, it wants to bill for the installation, the tuning and the maintenance too.

\$4B

The deployment company launched by OpenAI, with Bain, McKinsey and Capgemini as shareholders.

**≈ 150
engineers**

Picked up in one stroke through an acquisition, to come settle inside your walls.

1 billion

Profiles in LinkedIn: data as the moat, and prices climbing every year.

Anthropic announced its own deployment joint venture, backed by major financiers, openly copying the model that made Palantir. They're all following the same trail, because it leads to the same place: the inside of your company.

Why they're after your payroll

Let's take Mensch's formula all the way. If intelligence is now manufactured elsewhere, from electricity, then the value your teams produced is shifting too, toward those who manufacture that intelligence. Mensch said it before the French National Assembly with a chilling figure: a potential deficit of one trillion euros per year for Europe, the equivalent of ten percent of its payroll.

That is what's really at stake. These alliances aren't trying to replace your management software. They're trying to capture the share of value that until now flowed through salaries. When someone installs a system that does the work of ten people, the system's invoice replaces ten lines of payroll.

Consulting no longer sells a recommendation. It sells the conversion of your payroll into a subscription.

And this is where everything plays out: on price. Because people are willing to pay far more for a service that replaces labor than for mere software. Silicon Valley has theorized it perfectly. In a talk I strongly encourage you to listen to, bluntly titled *Software is Eating Labor*, a partner at the Andreessen Horowitz fund explains how software pricing has moved from price per user to price per outcome. His thesis is limpid and chilling: software first changed the way we store information, and its next act is to change the very nature of the economy, capturing trillions of dollars along the way. Translation: those trillions are today's salaries. When you bill for an outcome and not for a tool, you can charge almost as much as the salary you eliminate, while still staying cheaper than it. That's the heart of the model that's coming.

The lock-in trap

And this is where the reassuring reflex snaps shut. When a team spends six months integrating a system deep into your data, your processes and your compliance, that system becomes load-bearing infrastructure. You can no longer rip it out without bringing everything down. You then depend, for maintenance and every evolution, on whoever installed it. The trade press has a phrase for this: becoming hostage to endless consulting costs.

Let's be fair, the lock-in isn't a done deal. Consulting remains a multi-vendor market, and companies want to integrate AI with everything they already have, which no lab can do alone. The attempt is real; its success isn't yet. That's precisely why the window is open now.

The textbook case I know by heart: LinkedIn

Want to see what total, already-installed dependence looks like? Look at recruiting, it's my trade. LinkedIn holds the structured data of more than a billion people. No competitor will ever beat it. Even a system five times better would lose, for a very simple reason: a billion people are not going to retype their career history by hand into a new tool, and LinkedIn won't let go of them. The data is the moat, and it's uncrossable.

The result, you know it if you buy those licenses. Every year, at renegotiation time, the price goes up. It's a cash cow, because the dependence is total. And the supposedly magical solutions that present themselves as alternatives often live only off LinkedIn data retrieved through third-party vendors, by means whose legality and compliance could raise questions. One thing is certain: prices keep climbing. It's painful, but the dependence is locked in. Remember the mechanism, because it's going to repeat itself, identically, with AI.

The economic lock, and the arbitrage to come

Because here's what awaits you. Once a company's systems are juiced up on one particular vendor's AI, getting out costs so much that nobody does it. And I guarantee you one thing: the price of the token is going to shoot up. The giants haven't all found their footing yet, they're trading blow for blow. Anthropic has just reached profitability and raised billions once more, while OpenAI, despite enormous revenue, is still losing around fourteen billion dollars for the year and doesn't expect to break even before 2030. As long as the sector's business model isn't stabilized, their interest is crystal clear: make you dependent, then run up the bill.

And here's the trap, the real one. The gains are such that doing without will be impossible anyway. You will arbitrate, line by line, between a salary paid to a human and a token invoice. And even when the invoice goes up, it will remain cheaper than a salary, and it will produce more intelligence. It's mechanical, it's cold, and that's why you have to see it coming.

Hence one final vigilance, which is the heart of survival. The more you delegate to these systems, the more speed you gain, but the more skill you risk losing. And the day you've delegated everything without keeping anything, you'll be exactly where they want you: unable to do without them, and unable to judge what they're billing you. Delegate, yes. Amputate yourself, never.

What not to do, and what to do

I'm not telling you to ban the firms. They have real skills, and depriving yourself of all help would be absurd. I'm telling you not to hide behind them. Not to ask them what to do while signing with your eyes closed, letting them bill you millions and install their systems everywhere.

Challenge them. Ask which model they're installing, and why that one. Ask who owns what in the arrangement. Ask what happens the day you want to switch. Keep your hands on your data and on the choice of your tools. And above all, don't be fooled by the calendar: things are moving fast precisely because it's in their interest that you sign before you've understood.

The real protection isn't finding the right vendor. It's leveling up enough to no longer be at their mercy. Being able to judge, arbitrate, and decide for yourself. The exact opposite of the dependence they're trying to sell you.

Taking back control — we'll get there, that's the whole point of the end of this guide. But before knowing what to do, you have to follow the money to the end, and look at who really captures the value. That's the next chapter.

● SOURCES

-
- 01 MAY 2026 · OpenAI Deployment Company, launched in May 2026, four billion dollars, investors including Bain, McKinsey and Capgemini. OpenAI (openai.com), CIO Dive (ciodive.com).
-
- 02 MAY 2026 · Acquisition to bring in about 150 deployment engineers, capturing the implementation margin on top of the model margin. Digital Applied (digitalapplied.com).
-
- 03 2026 · Anthropic's deployment joint venture, backed by major financiers, on the Palantir model. MindStudio (mindstudio.ai).
-
- 04 2026 · The lock-in, load-bearing infrastructure and hostage to endless consulting costs. CIO (cio.com).
-
- 05 2026 · Nuance: consulting remains multi-vendor and companies want to integrate AI with their existing systems. AI Business (aibusiness.com).
-
- 06 MAY 2026 · Arthur Mensch before the French National Assembly, AI as the conversion of electricity into intelligence, potential deficit of one trillion euros per year, ten percent of Europe's payroll. Revue Conflits (revueconflits.com), Alsace IA (alsace.ai).
-
- 07 2025 · Alex Rampell (Andreessen Horowitz), *Software is Eating Labor*, a16z LP Summit 2025 (software pricing moves from per-seat to per-outcome, software is poised to capture trillions of dollars of economic value, accelerated by AI agents). a16z.com
-
- 08 MAY 2026 · Contrasting profitability among the labs. Anthropic on track for its first profitable quarter and close to a raise of about 30 billion dollars valuing the company around 900 billion, while OpenAI projects a loss of about 14 billion for 2026 and breakeven pushed back to 2030, with about 95% of ChatGPT users on the free tier. Tech Startups (techstartups.com), European Business Magazine (europeanbusinessmagazine.com).
-
- 09 2026 · On LinkedIn dependence and data as a barrier, as well as continually rising prices and the compliance gray zone of certain third-party data vendors: the author's field observation, cross-checked with public analysis of the sourcing ecosystem. To be presented as professional observation, without naming any player or alleging illegality.

Who captures the value, and who picks up the tab



Up to this point, this guide has talked about your job, your team, your company. This chapter follows the money. Everything before this described the same thing seen up close: value leaving work to settle somewhere else. Which leaves the question almost nobody really asks. Where does that value go, and who ends up paying the tab.

Intelligence is worth nothing anymore; the machine that produces it is worth everything

Take the formula from Mistral's CEO again — it has come up several times here. Intelligence, from now on, is electricity converted into tokens. It can be manufactured in near-unlimited quantities, and its price is collapsing, fifty to seventy percent a year on models, as we saw in chapter four. The service is becoming commonplace. In the strict sense, it becomes a commodity, a good without scarcity that nobody manages to sell at a premium anymore.

But look at the other end of the chain. If the service is worth nothing anymore, the means of producing it is following exactly the opposite path. Physical computing power — the processors, the data centers, the electricity that feeds them — is becoming the scarce resource. And a scarce resource, when the whole world starts wanting it, doesn't stay a mere expense line for long.

The service becomes free. The means of producing it becomes a treasure. All the wealth moves from the first to the second.

Compute just hit the trading floor

Here is the fact, and it's brand new. On May 12, 2026, the world's largest derivatives marketplace, CME, and Silicon Data, the company that publishes the first daily indices of GPU rental prices, announced the launch of the first futures market on computing power. It isn't open yet — it's waiting for the regulator's green light — but the direction is set. The CME's chairman summed it up in one sentence: compute is the new oil of the century. The founder of the fund behind the index went further: he sees it as the biggest commodity of the world to come.

Take the measure of what that means, because this is not a story for financiers. When a resource goes from being an operating expense to an asset that can be priced, hedged and traded, the same thing always happens, in the same order. Oil lived through it in 1983. From that moment on, those who held the physical resource stopped being mere suppliers. Capital started flowing toward them, their cost of financing collapsed, and a new aristocracy of rentiers was born — those who already owned the wells before the market turned them into gold.

May 12, 2026

The first futures market on computing power, announced by CME and Silicon Data.

\$50B

The cost of a single gigawatt of computing capacity. The entry ticket to a new kind of rent.

€1T/year

The deficit this shift could carve out for Europe — ten percent of its payroll.

For you, reading this guide from an HR department, the lesson isn't financial, it's brutal. The machinery being put in place has one single effect: guaranteeing that money will keep flowing into AI, no matter what. Don't bet for a second on a slowdown. Finance just built the pipe that ensures acceleration. So the pressure on jobs and skills won't slow down either.

What Mensch came to say, at the scale of a continent

That is exactly what Arthur Mensch, Mistral's CEO, came to explain before the French National Assembly's commission of inquiry on May 12, 2026. Go watch it — it's freely available, the link is at the bottom of the chapter. His thesis: AI is not software, it's infrastructure; it turns electricity into intelligence, and whoever controls that conversion captures a share of the value produced by everyone else. He put a number on the risk for Europe: a possible deficit on the order of a trillion euros a year, the equivalent of ten percent of the continent's payroll, flowing to those who manufacture that intelligence elsewhere.

And he stated the entry ticket. Building one gigawatt of computing capacity costs about fifty billion dollars. While American players are already committing trillions to corner the physical resources — processors and energy — Europe watches, for want of capital markets up to the task. The Draghi report had already traced the slope: Europe's share of global tech revenue fell from twenty-two to eighteen percent in ten years, while the United States' climbed from thirty to thirty-eight. Mensch puts a word on what comes next if nothing changes within two years: vassal states.

I'm not making this about flags. The question is not who wins the race. The question is that value is switching sides, and that movement is global. Europe simply sits on the wrong side of the conversion.

We had ideas, for want of oil

There's a slogan France grew up with, born of the oil shock. We have no oil, but we have ideas. Gray matter versus raw materials, the pride of a country that made up for the missing resource with ingenuity. That formula has just turned against itself. If intelligence is nothing more than transformed electricity, then ideas are no longer the immaterial, free resource we used to set against crude. They have become crude. A raw material extracted with energy, capital and silicon.

And here's the surprise: France has the new oil. Its low-carbon electron — nuclear, partly in surplus — is exactly the input needed to feed data centers. France isn't going from we have ideas to nothing. It's going to energy, and talent. Two cards out of three.

That leaves the third, and it's the whole game. Capital, and the place where the rent gets captured. Oil showed it: the people who get richest aren't always the ones pumping the crude, they're the ones trading it. The futures market on compute is opening in Chicago, not Paris. The risk for France, then, isn't running out of ideas. It's supplying the electron while the value flees elsewhere — into trading, into models, into the software layer. Being the electric emirate of a system it doesn't own. Rich in input, poor in rent.

Getting out of that trap comes down to three moves, and none of them is out of reach. First, convert the electron into intelligence here, on the soil where it's produced, instead of exporting it raw and then buying back, at full price, the intelligence extracted from it. Because that is exactly what's in the works — chapter eight showed it. The value goes off to be manufactured elsewhere, then comes back to be sold to European companies through the contracts the consulting firms have stitched up with the giants. We supply the material, we buy back the finished product. Losing twice.

The second move is more technical, and decisive. Never tie yourself to a single model or a single vendor. Build your use cases so you can swap engines without breaking everything; keep the process portable from one shop to the next. Perfect portability doesn't exist — it has a cost, and it caps the power a little. But being able to leave at a reasonable cost is precisely what the hostages of the lock-in described earlier were missing. It's the only insurance against the bill that keeps climbing.

The third: give vendors rooted here a real chance. Not out of patriotism — out of arithmetic. A vendor subject to European law, taxation and labor law reinjects part of the value into the very ecosystem that funds our social model. The one capturing the rent in Chicago does not. The point isn't to shut the others out, but to stop disqualifying European solutions from tenders by default just because they don't come from Silicon Valley. To let them take the lead when they can. It's less a posture than a long-term calculation.

The question nobody asks: who will pay for social protection

Now, the blind spot. The real one. The one I haven't heard mentioned anywhere — not from politicians, not at conferences.

In our economies, work doesn't just feed the person doing it. It funds social protection. Contributions levied on wages pay for pensions, healthcare, unemployment, long-term care. Our entire social model rests on a base, and that base is work.

Now, what has this guide described, chapter after chapter? Value is leaving work. It's migrating to capital — more precisely to an asset, computing power, that is being financialized as we speak, that will be held by a handful of players, and lightly taxed. Put the two ends together. The base that funds our social protection is draining away at the precise moment wealth is concentrating in something that isn't work, and that we don't know how — or don't want — to tax for it.

And this isn't a 2040 problem. Remember the tempo, from chapter six. Job cuts are immediate; new job creation is slow. Between the two opens a gap of several years. It's precisely during that gap that the transition would need to be funded — supporting people, paying for retraining. At the very moment, in other words, when the base starts to give way.

Work funds social protection. And work is what's being pulled out of the equation. Nobody has connected the two out loud yet.

Don't tell me the debate exists. It surfaced once, then faded. The robot tax, championed in 2017 by Benoît Hamon in France and floated the same year by Bill Gates, remained a campaign curiosity — theoretical, with no follow-up. Here and there, we have begun decoupling social financing from wages alone — the CSG, the so-called social VAT. But at the margins, never at the scale of the shock that's coming, nor on the factor that now captures the value. The debate isn't settled; it isn't even open. Because to open it, you would first have to state the finding. And the finding hasn't been made.

That's where the figure from Jérémy Lamri, the founder of Tomorrow Theory, takes on its full meaning. I heard him advance an idea at a conference that stuck with me: past a certain level of unemployment, something like fifteen percent, it's no longer a business problem — the whole system seizes up. Nobody can vouch for the exact threshold. The order of magnitude is enough. If employment contracts while the social base gives way, we're no longer talking about optimization, we're talking about a country's equilibrium. Those are the questions that need asking. Not in ten years. Now.

And me — what can I do

That's the question you're left with, after a chapter about exchanges, continents and billions.

Let me be honest about what I'd like, and about what I see. What I'd like is for everyone to wake up and really look. For nobody to stop at I typed three prompts into ChatGPT, it's cute, but it makes up half of what it says. Because what this guide has laid out — the curve, the equation, the jobs that are mutating, the value switching sides — is much, much bigger than that.

What I see is more down-to-earth. Until you've experimented yourself, touched the impact with your own fingers, you don't feel concerned. And you don't have the bandwidth to get started, because you're drowning in day-to-day management that doesn't always deliver the most value. You can't project yourself onto something you haven't lived through and don't have an hour to look at. I don't blame anyone for it. That's the real state of HR departments, and you have to start from there, not from where you'd like them to be.


Add a constraint that is no small detail. Your data is among the most sensitive there is — your candidates', your employees'. You cannot hand it over to just any tool plugged into just any cloud. You need to be able to trust, and trust can't be decreed.

That is exactly the observation my answer was born from. Not a cathedral, no ambition to replace your big systems. The opposite. Small, simple things, built with knowledge of the trade, that remove a chore and give you back an hour. That's the first step, and it's the whole point of the final chapter.

● SOURCES

-
- 01 **MAY 2026** · Hearing of Arthur Mensch, cofounder of Mistral, and Audrey Herblin-Stoop, before the French National Assembly's commission of inquiry into the structural dependencies and systemic vulnerabilities of the digital sector, May 12, 2026 (AI as infrastructure, converting electricity into intelligence, potential deficit of a trillion euros a year, i.e. ten percent of European payroll, one gigawatt of compute at roughly fifty billion dollars, trillions committed on the American side, risk of vassal states within two years). Official video, LCP-Assemblée nationale channel: [youtube.com](https://www.youtube.com). Revue Conflits (revueconflits.com), Alsace IA (alsace.ai).
-
- 02 **MAY 2026** · CME Group and Silicon Data, backed by market maker DRW, announce on May 12, 2026 the launch of the first futures market on computing power, expected later in the year subject to regulatory approval, indexed on the first daily benchmarks of GPU rental prices. Compute is described there as the new oil and the world's biggest commodity to come. CME Group (cmegroup.com), PR Newswire (prnewswire.com), Hedgeweek (hedgeweek.com).
-
- 03 **SEPTEMBER 2024** · Draghi report: Europe's share of global tech revenue down from 22% to 18% between 2013 and 2023, the United States up from 30% to 38%. Cloud News (cloudnews.tech).
-
- 04 **2025** · Unemployment threshold past which the system seizes up, around fifteen percent: remarks by Jérémy Lamri (Tomorrow Theory) heard at a conference, reported in substance, without a verbatim quote for lack of a public recording. Profile: dunod.com.
-
- 05 **2017** · Robot tax, proposed by Benoît Hamon during the 2017 French presidential campaign, an idea floated the same year by Bill Gates. To be cross-checked against primary sources at publication time (Hamon 2017 platform, Bill Gates interview with Quartz, February 2017).
-
- 06 Social protection financing based mainly on payroll contributions in France and Switzerland, with a partial shift toward tax-based funding already under way (CSG, social VAT). A framing observation to be backed by an institutional source at publication time (for instance DREES in France, OFS in Switzerland).

Survival means going local



Here we are at the end. Nine chapters spent looking squarely at what's playing out. The curve that won't wait for you, your ERP's wall turned into a door, the new value equation, the jobs that are mutating, the alliances aimed at your payroll, the value switching sides, and the social tab nobody wants to see. One question remains, the only one that really counts. Now what do we do?

Local, or responsibility regained

My answer fits in one word, and its technical modesty will surprise you: local.


Running AI locally — that is, on your own machines, with your own data that never leaves your premises — is not an engineer's whim. It's a position. When a model runs on your premises, your candidates' and your employees' data stays with you. So does the judgment. You're not renting your intelligence from a vendor who, somewhere else, is recording its own employees' every move the better to replace them. You stay in control.

And what makes this possible today is new

For a long time, local was wishful thinking — too slow, too weak. That changed in a matter of months, and I'll spare you the technical details, but hold on to two movements. First, new methods are sharply accelerating models on ordinary machines — multi-token prediction, for instance, which delivers a speed gain of one and a half to two times through software optimization alone, without changing hardware, or the quantization techniques that slim models down. If you want to dig in, the sources are at the bottom of the chapter. Second, and this is the most important part, open models are no longer years behind proprietary ones. They're months behind.

That's not an impression, it's measured. The Epoch AI institute calculated the capability gap between the best open models and the best closed models. Since January 2026, that gap has averaged just four months.

Open models are only months behind the best closed models



Source: Epoch AI, *Epoch Capabilities Index*, May 2026. Chart reconstructed under a Creative Commons BY license.

Take the measure of the opportunity. Models you can host yourself, with nothing leaving the building, are catching up with the most advanced proprietary systems. Mistral in France, or Chinese labs like Qwen and DeepSeek, are releasing lightweight versions that already deliver serious results on concrete tasks. Without compromising the sovereignty of your data, without sending anything outside, and above all without chaining yourself to a single vendor. Because the day the models evolve, and that's a certainty, your systems can evolve too — at your pace, not theirs.

What I do with it, concretely

That is exactly the choice I made. Not for large corporations — I'm not an IT engineer and they have their own teams. But for my community: HR people in French-speaking Switzerland, and the recruiting world more broadly. Giving them autonomy without compromising themselves with a vendor, keeping the most sensitive data encapsulated in a local system, and freeing people from the tasks that bring no value to the trade. Or, better yet, letting them do what a well-trained human knows how to do, but faster and at scale. Doing better, in short. And that's good for everyone.

Remember the formula from Mistral's CEO, met earlier in this guide. AI is converting electricity into intelligence. The real question is where that conversion happens, and who profits from it. At the scale of a continent, that's sovereignty. At the scale of your HR department, it's exactly the same fight — smaller, and every bit as decisive.

Keep the data, keep the judgment, keep control.

Not to add a layer of admin — to remove one. Because it's by freeing themselves from valueless tasks that HR people will finally be able to take up the strategic subjects that are rightfully theirs. I'm not telling you this in theory. I do it, and it works.

Above all, don't panic

Now, a warning, and it matters as much as everything else. Becoming aware of this shift does not mean rushing. Quite the opposite.

Remember Klarna. That company, whose CEO is one of the most AI-enthusiastic people on the planet, charged in headfirst. Hiring freeze, headcount cut from five thousand five hundred to three thousand four hundred, a chatbot presented as the equivalent of seven hundred agents. Six months later, quality was collapsing, customers were fleeing, and the company was rehiring humans. The CEO's admission became famous: we focused too much on efficiency and cost, the result was lower quality, and that's not sustainable.

5,500 → 3,400

Klarna slashed its workforce in the name of AI... before rehiring humans.

55%

Of companies that rushed to replace humans with AI regret it.

2027

The year by which half of those that cut customer service will have to rehire (Gartner).

This is not an isolated case. More than half of the companies that rushed to replace humans with AI regret it, and one research firm predicts that by 2027, half of those that cut into customer service will have to rehire.

The lesson is crystal clear, and it's the first rule of survival. Don't fire anyone right away. The cut-costs-first reflex is the most expensive of all, because we always forget the price of undoing what we broke.

Now, here is what to do

No miracle solution. Four steps, in order.

First, become aware. Reread these articles, open the sources, verify everything. Take the facts, feed them into an AI if you feel like it, discuss them with it, ask your questions. Learn. It's free, and it's the only investment that will never turn against you.

Second, alert your leadership. This shift can't be managed from a corner of the HR department. Someone at the top has to take it on. Not necessarily a board member, but clearly someone in senior management who grabs the subject with both hands and decides to understand. Look at Ivalua: the founder left his CEO seat to become Chief AI Officer, a role created just for this, explaining that you need a global view of the whole company and its processes to make the change happen. When a founder switches chairs for that, the signal is clear.

Third, get your hands dirty. Don't be afraid. You don't know how? Nobody knows how — it's new for everyone. So start in a sandbox, run a hackathon as we saw, tinker with things, for yourself or for your company. Your asset is knowledge of the trade, and it's the most precious one, because the technical barrier between a business problem and its solution has almost disappeared. You may need to surround yourself with people who aren't afraid to get their hands into the technology, but it's not as complicated as it used to be. What counts is going in.

And to go in, there is one condition — just one, but it's absolute. You have to free up the time. As long as you're bogged down in the time-consuming, valueless tasks of your job, you'll never have an hour to think, analyze, step back. It's the snake biting its own tail, and that's the whole point: freeing yourself from those tasks is precisely what will give you the time to learn how to free yourself from them. That's why the very first move, lifting the administrative weight, isn't a comfort. It's the key to everything else.

Fourth, once you've touched the tools with your own hands, dare an audit. With me or with someone else, it doesn't matter. And it will be far more useful once you've experimented, because you'll know what to look for. The goal isn't even to find solutions right away. It's to become aware of your weak points, to gauge how fast things are moving, and to understand where AI could bring something to your shop. You can perfectly well detach a few people to go explore. Because the real question isn't do you want to go. It's: if you don't, your competitors will. And then what happens?

And all of this has a name

These four steps — become aware, give the subject an owner, get your hands dirty, dare the audit — are not a fad. They're a discipline. In recruiting it has a name, RecOps, and it's the trade I practice. But never mind the word. RecOps for recruiting, TalentOps, HROps, Ops of whatever you like, the moment you're looking at systems and processes. Or Forward Deployed Engineer, if it's purely AI applied to company processes. The title doesn't matter at all. What matters is the movement: stop submitting to the tool, and get back to driving it.

And this is where I add a string to that bow. RecOps sets the course — the right processes, the right goals. But between the course and the daily life of a bogged-down team, a bridge is almost always missing. That bridge is small custom tools. Not bloated monsters: simple objects, cut for a specific trade and a specific problem, that remove the chore eating up the time. I build them as a craftsman, with a recruiter's command of the trade, and with one rule that never moves: rigor on risks and on data first, power second. Even if it means offering something smaller than the giants do. The open models we run locally are improving so fast, as we saw above, that the gap closes month after month. It's my way of reconciling the two — what we should do, and what we can actually do, starting Monday morning.

And the real skill, the rare one, isn't technical. It's knowing what you want to do. Understanding well enough what these tools make possible to walk into a company, spot what can be transformed, and turn that into a roadmap. That's where the opportunity lies, and it's immense, for whoever wants to seize it.

Let's be clear: as of mid-2026, some things aren't feasible yet. But at the speed everything is moving, what looks crazy today will be normal, and simple to set up, in 2027. The only mistake would be to wait until it's easy before starting to understand.

That's why this guide is called a survival guide. Not because the ending is written. On the contrary. Because nothing is settled — provided you look, understand, and take back control while there's still time.

I could have joined the race and helped speed up the machine. I chose the opposite: helping people keep their footing. That's all this guide will have tried to do. The rest is up to you.

I'm a sourcer. If you look for me, you'll find me.

● SOURCES

- 01 **2026** · On accelerating local inference: multi-token prediction (MTP) delivering 1.5 to 2x gains through software optimization alone with no extra hardware, and quantization (including TurboQuant) to slim models down. Google, MTP drafters for Gemma 4 up to 3x faster ([blog.google](#)), llama.cpp tutorials and benchmarks ([datacamp.com](#)), overview of 2026 local hardware and open-model parity on everyday tasks ([presenc.ai](#)).
-
- 02 **MAY 2026** · Epoch AI, *Open models lag state-of-the-art closed models by 4 months*, May 29, 2026 (since January 2026, open models trail closed models by four months on average, i.e. eight ECI points): [epoch.ai](#). Chart reconstructed from the published data, under a Creative Commons BY license. Fast Company ([fastcompany.com](#)), Entrepreneur ([entrepreneur.com](#)), Digital Applied ([digitalapplied.com](#)).
-
- 03 **2026** · Orgvue and Forrester: 55% of companies that rushed to replace humans with AI regret it, via Linkifico ([linkifico.com](#)). Gartner: by 2027, half of those that cut customer service will have to rehire, via Vibe Graveyard ([vibegraveyard.ai](#)).
-
- 04 **JANUARY 2025** · Ivalua: the founder moves from CEO to Chief AI Officer. PRNewswire ([prnewswire.com](#)), Technology Magazine ([technologymagazine.com](#)).
-
- 05 **2026** · On local deployment and data staying home, to be set against the employee surveillance deployed outside the European Union, covered in chapter 8. EMARKETER ([emarketer.com](#)).

• ONE LAST THING

*I'm not asking you to
believe me. I'm asking
you to look at the same
facts I did.*



— THE AUTHOR

Guillaume Alexandre

A recruiter for more than 15 years, founder of Gates Solutions, based in Le Brassus in the Vallée de Joux. A sourcer by trade, he has connected his own recruiting tools to AI, automated what could be automated, and works with software vendors to push those tools beyond what they can do natively.

His conviction: helping companies, professions and regions take AI into their own hands rather than endure it. Keeping the data, the decisions and the value close to the people who produce them. This guide is the first stone.

"I'm a sourcer. If you look for me, you'll find me."

• STAY IN TOUCH



FOLLOW THE AUTHOR

Guillaume Alexandre

LinkedIn →



THE WEBSITE

gates-solutions.com

Visit →



ALSO PUBLISHED AS A BOOK

Buy the printed edition

French edition · Amazon →



THE COMPLETE GUIDE

Download as PDF

Direct download →



WORTH PASSING ALONG

Send by email

Share this guide →



LE BRASSUS · VAUD · JUNE 2026 · "EVERYTHING IS SOURCED"