

TOLD THROUGH THE NUMBERS ON THE TABLE

The Control Gap

AI's adoption curve, by the numbers — what one room full of technology leaders revealed about the state of AI.

THE STATE OF PLAY

On 2 July 2026, senior technology and digital leaders gathered at The Plectrum in Ghent for the Dropsolid × Beltug event to ask a hard question: are organisations actually in control of the AI they've adopted?

The conversation ranged across governance, cost, security and sovereignty — but the numbers on the table told the story better than any single argument could. Here they are, in order.

67%

of people are already using AI in their day-to-day work

18%

of organisations have an AI governance framework in place

That's not a gap. That's a **chasm** — and it's the current state, not a future risk.

Two out of every three employees already rely on AI to get their work done — often on their own initiative, often without anyone in leadership knowing which tools, which data, or which risks are involved. Fewer than one in five organisations has built the guardrails to manage that reality.

The other 82% aren't behind schedule — they don't have a schedule. Adoption doesn't wait for governance to catch up; it compounds. Every unsanctioned tool and every "quick win" widens the space between what people do with AI and what the organisation can see, explain, or control.

- Adoption is a rounding error away from universal. Governance is still the exception.

WHAT'S AHEAD

01 The Reckoning — why the risk, the failure rate, and the real cost of AI are bigger than most budgets assume.

02 The Exposure & the Way Through — Europe's sovereignty gap, and the proof that getting this right is possible.

PART ONE

The Reckoning

Four numbers make the risk impossible to ignore. AI doesn't invent new weaknesses — it finds the old ones faster, and the room's own data proved it.

Risk is already
showing up

50%+

have already had a negative AI-related incident in the past year

3×

increase in external cyberattacks over the past two years

Most projects don't
survive contact
with reality

The technology is rarely the reason AI projects die. The organisation around it is.

95%

of AI projects never make it past the pilot stage

10/20/70

the tech / change / people split effort should follow — most do the inverse

The bill is bigger
than the invoice

Licences measure spend, not value — and today's subsidised token prices were never going to last.

3–5×

what adoption, change and tuning cost on top of the licence fee

\$100B

estimated spend on the infrastructure behind today's AI — someone must recover it

PART TWO

The Exposure, and the Way Through

Europe's dependency — and the proof that control is possible.



of Europe's technology stack is non-European. On infrastructure, its presence is close to negligible.

The cloud it runs on, the models it calls, the chips that train them, the platforms that orchestrate them — for the overwhelming majority of European organisations, none of it is theirs to control. It's rented, from a small handful of providers almost all outside Europe. That's not a footnote to the AI story. It's the foundation everything else sits on.

This would matter little if dependency were theoretical. It isn't. A brief regulatory disruption to frontier-model access, during the very event this report is drawn from, made the risk concrete in real time. The lesson wasn't that any vendor is untrustworthy — it was that discovering you have no alternative during a disruption is far worse than discovering it on a calm afternoon beforehand.

So the fix isn't isolation, and never was. No organisation — and no country — can build everything itself. The realistic goal is optionality: knowing where your dependencies sit, deciding which you're comfortable carrying, and having a genuine alternative for the ones you're not. An organisation with no Plan B has no leverage. One with a credible Plan B does — whether it ever activates it or not.

Sovereignty was never about isolation. It's about **optionality** — a credible Plan B changes your negotiating position, even if it's never activated.

THE PROOF

None of this is unsolvable. The proof is already on the record.

93%

reduction in manual processing after applying AI to a well-scoped problem — thousands of weekly job applications at the EU Personnel Selection Office.

These aren't pilots, demos or projections. They're measured outcomes from organisations that did the unglamorous work first: define the problem precisely, decide upfront what success looks like, redesign the process around the answer, and keep someone accountable for the result. Neither came from adopting AI faster than everyone else — they came from adopting it more deliberately.

- The technology clearly works — when the problem is well-defined, the process is redesigned around it, and someone stays accountable for the result.

95%

of the time, AI-powered penetration testing found more issues than a traditional human test — across 1,000+ tests in ten months.

That's the real difference between the exposed majority and the handful who aren't. Every organisation here had access to comparable tools. The difference was which ones treated control as a design decision made *before* deployment, rather than a cleanup job attempted after.

THE THREAD

Control is not something organisations lose to AI by accident. It's something they hand over by failing to ask hard questions early enough — about dependencies, about what success actually means, about who's accountable when something goes wrong. The organisations willing to ask them now are the ones still standing — and still in control — when the next disruption arrives.