

DIRTY DATA EXPERT

Structure Before Intelligence™

THE €825 MILLION LESSON

Executive Brief

Gary Segal

www.dirtydataexpert.com

gary@dirtydataexpert.com

Executive summary

In August 2026 the Dutch data protection authority fined Uber €824,990,000 because software, not people, decided when drivers lost access to their income. It is the second-largest GDPR fine ever, and it had nothing to do with a data leak or a hack.

The regulator did not object to Uber using automation. It objected to what was missing around it: no human review before a consequential decision, and no explanation to the people affected. That is a structure failure, and it is exactly the gap Structure Before Intelligence™ exists to close.

Uber's system was rule-based automation, not a modern AI agent. That makes the case more relevant, not less. Agentic AI adds autonomy, speed and scale to the same pattern, so any organisation deploying agents is now on notice for what regulators expect around automated decisions.

The fine was not for using an algorithm. It was for letting one decide with no structure around it.

What happened

The Dutch Autoriteit Persoonsgegevens (AP) found that Uber breached the GDPR's ban on fully automated decisions with significant effects, and drivers' right to be told about them. The decision is dated 17 August 2026 and was confirmed publicly on 21 August.

Item	Detail
Fine	€824,990,000, second only to Meta's €1.2bn (2023)
What the software did	Temporarily deactivated drivers suspected of fraud; permanently deactivated drivers with persistently low customer ratings
What was missing	No human assessment before the decision took effect; drivers not adequately told the decision was automated
Period	Reported as 2018–2022 by law firms, 2020–2022 by Dutch press
Impact on people	Deactivated drivers lost their income for the duration
How it started	Complaint by French drivers via France's CNIL; handled in the Netherlands, where Uber's European HQ sits
History	Uber's fourth AP fine, after €600,000 (2018), €10m (2023) and €290m (2024)
Status	Under appeal

Uber's position. Uber calls the fine disproportionate and says permanent deactivations were never fully automated. It says only 126 drivers in Europe were permanently deactivated for low ratings in 2021. It also says its current process includes human review and a way for drivers to appeal.

That last point matters for leaders: the fix Uber now cites is the structure the regulator says was missing at the time.

Why this is a structure failure, not a technology failure

Nothing in the decision says Uber's software malfunctioned. The rules did what they were built to do. What failed was everything around them.

- The inputs were never questioned. Customer ratings and fraud signals were treated as facts precise enough to end someone's income. Ratings are noisy, subjective data. Nobody decided how reliable an input must be before it can trigger a high-stakes action.
- No one owned the decision. Once the threshold was crossed, the system acted. There was no named point where a person reviewed, approved or could stop it.
- The people affected could not see in. Drivers were not adequately told a machine had made the call, so they could not challenge it on fair terms.
- Scale hid the risk. Uber argues few drivers were affected. The regulator still fined it €825m, which signals that the missing safeguard itself is the offence, whatever the number harmed.

This is the Structure Before Intelligence™ argument in its plainest form. You cannot bolt decision-making automation onto a structure that was never designed to carry it. The cost of that gap now has a public price tag.

Where the structure broke: the four layers

Every layer of the Structure Before Intelligence™ house shows a gap in this case. The Structure layer, governance and policy, is the one the regulator actually fined.

Layer	What it should provide	What the Uber case shows	The question to ask
Foundation (data)	Inputs reliable enough for the decision they drive	Customer ratings and fraud signals used as grounds to cut off income	Is this data good enough to justify this consequence?
Structure (governance)	Rules on which decisions machines may take alone	No human review before a significant decision; no disclosure to drivers	Which decisions must a person sign off, and who is accountable?
Walls (process)	Workflows with checkpoints, escalation and appeal	Decisions took effect with no review step or route to challenge	Where is the human checkpoint, and how does someone appeal?
Roof (leadership)	Sponsorship that funds oversight, not just automation	Four AP fines since 2018 suggest governance did not keep pace with growth	Does the board see automated decisions as a risk it owns?

Uber's own remedy, human review plus an appeal route, sits in the Structure and Walls layers. That is where the money should have gone first.

Why AI agents make this risk bigger

Uber's system was simple threshold automation. An AI agent is the same pattern with more autonomy: it reads data, decides and acts across systems without waiting for a person. Every gap in the Uber case gets wider.

- More decisions, faster. An agent can suspend accounts, reorder stock, reject invoices or reprice products thousands of times a day. A missing checkpoint fails at that volume too.
- Harder to explain. Uber's rules could at least be written down. An agent's reasoning can shift between runs, which makes the GDPR duty to inform affected people harder, not easier.
- Dirty data becomes action, not just a bad report. In supply chains, roughly eight in ten organisations already cite data limitations as the main barrier to scaling agentic AI. An agent acting on poor master or inventory data automates the mess at speed.
- Oversight is not keeping up. In SANS's 2026 survey, 76% of organisations reported growth in non-human identities such as bots and agents, and only about four in ten use human approval for agent actions.
- The rules are converging. GDPR already covers automated decisions about people. The EU AI Act's enforcement powers went live on 2 August 2026, and employment and worker-management uses are classed as high-risk.

For FMCG, retail, manufacturing, logistics and financial services, the exposure is concrete: agents deciding on supplier onboarding, credit limits, customer accounts, staff scheduling or returns. Any of those that affect a person significantly sit in the same legal territory as Uber's driver deactivations.

Before an agent gets the authority to act, decide what it may never do alone, who reviews it, and how an affected person finds out and objects.

Questions for the leadership team

1. Which decisions in our business are already made by software, rules or AI with no human review, and which of them affect a customer, supplier or employee significantly? 2. For each one, how reliable is the data that triggers it, and has anyone tested it? 3. Who is the named owner accountable for each automated decision? 4. Can an affected person find out a machine decided, and challenge it? 5. Before we give any AI agent authority to act, which actions require human sign-off? 6. If a regulator asked for evidence of our safeguards tomorrow, what could we show them?

Work with Gary

Keynotes, executive workshops and consulting on AI readiness, data governance and Structure Before Intelligence™.

gary@dirtydataexpert.com · +27 68 645 1673 · www.dirtydataexpert.com

Sources

1. Reuters, via Euronext: Dutch regulator fines Uber \$966 million for automating driver suspensions
 2. NOS: Uber schond rechten chauffeurs (Dutch)
 3. Pearl Cohen: Dutch DPA fines Uber €825 million
 4. Baylos: Dutch DPA fines Uber €824.99 million
 5. Recording Law: GDPR fines tracker
 6. Autoriteit Persoonsgegevens announcement
 7. Supply Chain Research: Agentic AI in the supply chain
 8. Infosecurity Magazine: SANS survey on AI agents and non-human identities
 9. Volkov Law: EU AI Act enforcement is no longer theoretical