THE MOMENT YOUR TEAM IS NAVIGATING

AI moves fast. Regulators are catching up.

Every company running AI on sensitive data is in the same position: the models are already in production, but the privacy infrastructure wasn't built for them. The window to fix this before an enforcement action or breach is narrowing.

Three converging pressures — rising breach costs, tightening AI regulation, and audit scrutiny of automated systems — are making this a board-level issue in 2025.

BY THE NUMBERS

$4.88M
Average cost of a data breach in 2024 — up 10% from 2023
IBM Cost of a Data Breach 2024

$9.77M
Healthcare sector average — the highest of any industry for 14 consecutive years
IBM / HHS 2024

€1.2B
Total GDPR fines issued since 2018 — Meta alone: €1.3B in a single action
GDPR Enforcement Tracker

277 days
Average time to identify and contain a breach — most start in AI pipelines or logs
IBM 2024

WHAT MOST TEAMS GET WRONG

- **Treating PII redaction as a post-processing step.** By the time data reaches a redaction layer, it's already in model context, logs, and embeddings.
- **Relying on LLMs to "not remember" sensitive data.** Models are trained on what they see. Prompts are logged. Context windows persist.
- **Manual redaction workflows at scale.** Teams of 5–15 people reviewing outputs doesn't scale past tens of thousands of daily records — and creates its own access risk.
- **No audit trail for AI access.** Regulators now ask specifically: who accessed what data, when, for what purpose? Most teams can't answer this for AI systems.

WHAT THE RIGHT APPROACH LOOKS LIKE

01 **Enforce at the data layer, not the model layer**
Strip sensitive fields before they enter any system — not after. Privacy by design, not by policy.

02 **Minimum necessary as a technical constraint**
Each workflow should only see the fields it actually needs. Not a configuration suggestion — an enforced API rule.

03 **Consent and purpose recorded automatically**
Every access event tied to a declared purpose and subject record — ready for regulator requests without manual reconstruction.

04 **Breach containment baked into architecture**
When a downstream system is compromised, it should only expose redacted data. Breaches happen — their blast radius shouldn't include raw PII.