The Oldest AI Law Is Already Being Enforced: GDPR and the Automated Decision

Basil C. Puglisi, MPA
A Human-AI Collaboration

For the executives who deploy AI on EU personal data, the lawyers who defend them, and the privacy officers who answer for it. Verified as of June 13, 2026.

How AI Legal Exposure Travels

Most organizations treat the General Data Protection Regulation as a cookie-banner and consent problem, settled years ago and parked with the privacy team. That reading misses what the regulation has become. The GDPR, in force across the European Union since May 25, 2018, is the oldest law on the books that directly governs automated decisions about people, carrying forward a rule that traces to Article 15 of the 1995 Data Protection Directive, and in 2026 it is one of the most active. What matters is the active enforcement rather than the age. It reaches organizations established in the European Union wherever the processing happens. It can also reach organizations outside the Union when they process the personal data of people in the Union to offer them goods or services or to monitor their behavior. That means a company in New York that runs an AI screening tool on European applicants can fall inside its scope.

Two provisions carry the weight for anyone deploying AI. Article 22 gives a person the right not to be subject to a decision based solely on automated processing that produces legal or similarly significant effects, with narrow exceptions for contract necessity, explicit consent, or authorizing law. For the contract and consent routes, the controller must put safeguards in place, at least the right to obtain human intervention, to express a point of view, and to contest the decision. For the authorizing-law route, the safeguards are set by that law. Articles 13, 14, and 15 require that people be told when their data feeds automated decisions and profiling, and be given meaningful information about the logic involved. The penalties sit in Article 83, and they are tiered.

For breaches of core processing principles and data subject rights, the ceiling is the higher tier: up to 20 million euros or 4 percent of total worldwide annual turnover, whichever is greater. Lesser organizational failures sit at up to 10 million euros or 2 percent. These are maximums applied case by case, scaled to severity and turnover under the European Data Protection Board's fine-calculation guidelines, not automatic figures. The point for a deployer

is the exposure category: an AI tool that makes or heavily shapes decisions about people can land in the higher tier, where the cost is measured against global revenue.

This unit covers GDPR alone. The EU AI Act runs alongside it with its own obligations and its own penalties, and the two apply together rather than one replacing the other. Sector rules add further layers. Those regimes sit outside the scope of this article.

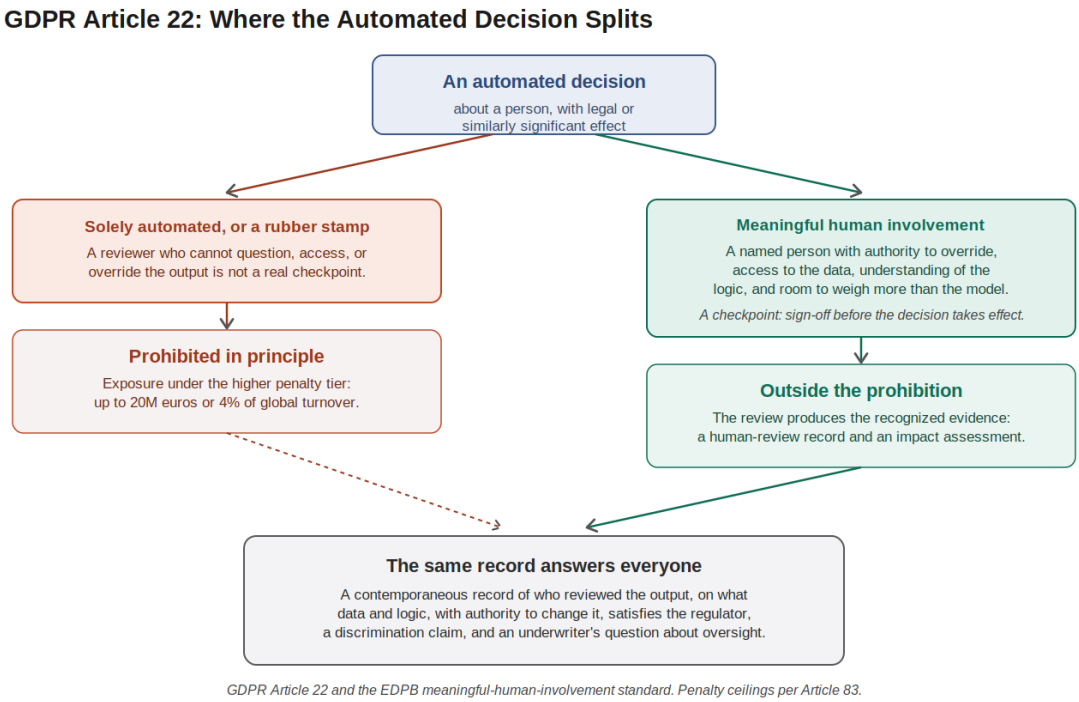

Figure 1. GDPR Article 22, where the automated decision splits: a solely automated decision or rubber stamp falls under the prohibition, while meaningful human involvement produces the recognized record.

Liability Already Tested in Court

The reason GDPR belongs at the front of any AI risk map is that its automated-decision rules have already been tested at the highest level, and the test landed on the exact point where AI governance lives.

In December 2023, the Court of Justice of the European Union decided the SCHUFA case, Case C-634/21 (judgment of December 7, 2023, ECLI:EU:C:2023:957). The court held that an automated credit score is itself an automated individual decision under Article 22 when a third party draws strongly on that score to grant or deny something to the person, such as a loan. The significance is the reach. A company cannot escape Article 22 by pointing to a

human who formally signs off, if the automated output is what effectively drove the result. The court looked at the whole decision-making situation, not the formal title of whoever held the pen.

That ruling captures the scorer. The test for the human step that follows comes from a separate source, the Article 29 Working Party guidance the European Data Protection Board later endorsed, and it is the heart of the matter. That guidance holds that human involvement in an automated decision must be substantive and capable of influencing the outcome. A rubber stamp does not count. For the human step to take a decision outside the solely-automated prohibition, the reviewer needs the authority to change or override the decision, access to all the relevant data, an understanding of the logic and criteria behind the automated output, and the ability to weigh information the system did not consider. A reviewer who clicks approve on a screen full of model outputs they cannot question is, in the regulator's reading, not a meaningful checkpoint at all.

Read plainly, the GDPR already requires what good AI governance describes. It asks for a named, empowered person at the decision who can actually change the answer, more than a human somewhere in the building, and it has asked for this since 2018.

The enforcement posture in 2026 makes the exposure concrete rather than theoretical. The European Data Protection Board's 2026 Coordinated Enforcement Framework targets transparency and information obligations under Articles 12, 13, and 14. It examines whether organizations give people clear, complete, and accessible notice, including notice that automated decision-making or profiling is being used where the GDPR requires that disclosure. Separately, AI processing is being folded into existing GDPR enforcement, with national authorities across the European Union pursuing AI-related actions that reach large-model training data and chatbot personal data, some of which remain contested or procedurally unsettled. The exposure is not waiting for a new AI-specific statute. It is being charged under a law that has been in force for years.

The Records the Regulation Already Names

The same provisions that create the exposure also name the evidence that answers it, and that evidence is the upside for an organization that produces it before a regulator or a claimant asks.

The first recognized artifact is the Data Protection Impact Assessment, required under Article 35 for processing likely to result in high risk to people, which automated decision-making and large-scale profiling routinely trigger. A completed, current impact assessment is the document a supervisory authority expects to see first, and a required assessment that is missing can itself support a finding, generally under the lower penalty tier of up to 10 million

euros or 2 percent of turnover, while breaches of the data subject rights in Articles 12 to 22 reach the higher tier.

The second is the human-review record. Where an organization relies on meaningful human involvement to keep a decision outside the Article 22 prohibition, the defensible position is a contemporaneous record showing who reviewed the automated output, what data and logic they had in front of them, and that they held the authority to change the result.

The third is the transparency record, the documented proof that clear, accessible notices under Articles 12, 13, and 14 were given where required, which is exactly what the 2026 enforcement action is examining.

These artifacts share a quality worth holding onto. They are recognized, auditable, and credited by the people who decide exposure, the supervisory authority, the claimant's lawyer, and increasingly the insurer pricing a privacy or technology policy. An organization that can produce them holds evidence that travels, because the same human-review record that satisfies Article 22 may also help answer a discrimination claim, a contract dispute, and an underwriter's question about oversight.

There is a second, more honest reward worth naming without overselling it. An organization that builds a real human checkpoint into its automated decisions, rather than a rubber stamp, tends to surface model errors and edge cases that a fully automated pipeline would have shipped. The benefit there is not only compliance. It is the quality and trust that come from a person who can catch what the system missed, and it shows people being kept in the loop rather than removed from it. That claim is modest and carrier-dependent, and it is offered as an observation rather than a guarantee.

Governance That Produces the Evidence

The mechanism that turns this exposure into those artifacts is a checkpoint, defined simply as a fixed point in the workflow where a named human must sign off before an automated decision takes effect or an output reaches the person it concerns. The recognized baselines for what good practice looks like are the established standards, the NIST AI Risk Management Framework and ISO/IEC 42001. Checkpoint-Based Governance, a method proposed by the author, is one approach to producing the specific evidence the GDPR calls for, offered as one way to do it rather than as a standard in its own right.

The connection to Article 22 is direct. The regulation's meaningful-human-involvement test asks for a reviewer with authority to override, access to the data, understanding of the logic, and room to consider more than the model saw. A checkpoint built to that test produces, as a byproduct, the human-review record that serves as evidence. The governance and the

documentation are the same act. The person governs the decision, and the trail of that governance is the artifact a regulator or court later reads.

The pattern is visible in a single regulation. The liability an organization must avoid is a four-percent-of-turnover finding for an unaccountable automated decision. The growth it wants to claim is a defensible and trusted decision process. Both are produced by the same move: putting a named human with real authority at the point where the automated decision is made, and keeping the record.

The Limits Worth Naming

A documented checkpoint is persuasive evidence, not a guarantee, and three limits deserve naming. First, a human-review record can cut both ways. If the record shows a reviewer who approved in seconds without the data or authority the regulator requires, it becomes evidence of a rubber stamp rather than a defense, so the checkpoint has to be real to help. Second, adopting a standard and then ignoring it can deepen exposure, because a documented process an organization failed to follow can support a "should have known" argument. Third, the meaningful-human-involvement test itself is not fully settled. Regulators and courts are still working out exactly how much human authority is enough, and proving from the outside that a human genuinely drove a decision is difficult. None of this removes the value of the record. It means the record must reflect real oversight, honestly kept.

The GDPR is the oldest law of its kind, and in 2026 it is among the most active, charging AI exposure under rules that have been in force since 2018. The evidence it rewards is already defined, the human-review record sits at the center of it, and the methods to produce that record are published and available.

Sources

Regulation (EU) 2016/679 (General Data Protection Regulation), Articles 13, 14, 15, 22, 35, and 83.

Directive 95/46/EC, Article 15 (the automated-decision rule GDPR Article 22 carries forward; the Directive was repealed when the GDPR took effect).

Court of Justice of the European Union, *SCHUFA Holding (Scoring), OQ v Land Hessen*, Case C-634/21, judgment of December 7, 2023, ECLI:EU:C:2023:957 (automated credit scoring as an automated individual decision under Article 22 where a third party draws strongly on the score).

Article 29 Working Party, Guidelines on automated individual decision-making and profiling for the purposes of Regulation 2016/679 (WP251rev.01), adopted February 6, 2018, endorsed by the European Data Protection Board on May 25, 2018 (the meaningful-human-involvement standard: substantive review with authority to override, access to data, understanding of the logic, and ability to consider additional information).

European Data Protection Board, CEF 2026 coordinated enforcement action on transparency and information obligations under the GDPR, Articles 12, 13, and 14 (topic selected October 14, 2025; launched March 19, 2026).

European Data Protection Board, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models.

European Data Protection Board, Guidelines 04/2022 on the calculation of administrative fines under the GDPR (severity-scaled, turnover-based maximums).

NIST AI Risk Management Framework (AI RMF 1.0) and ISO/IEC 42001:2023, recognized governance baselines.

Disclaimer

I am not a lawyer, and this article does not provide legal advice. This is thought research and governance analysis based on public sources, cited materials, and human-AI review. It is intended to help executives, practitioners, insurers, and governance teams think more clearly about AI risk, liability exposure, and documentation practices. Readers should not rely on this article as a legal opinion, compliance determination, or substitute for qualified counsel. Any organization facing a legal, regulatory, contractual, or insurance question should consult its own attorney, broker, or professional adviser before acting.

The Other AI: Audio Briefings on Augmented Intelligence and AI Governance
Spotify | Apple Podcasts | Amazon Music | YouTube Playlist

#AIassisted using HAIA Ecosystem