EU PRODUCT LIABILITY DIRECTIVE

(DIRECTIVE (EU) 2024/2853)

EXPOSURE TEST

INTRODUCTION

Most organisations still approach AI and software risk through the language of compliance: policies, controls, approvals and governance frameworks. Those things matter, but they are rarely the questions that arise when something has gone wrong. Once harm is alleged and a complaint, claim or legal challenge follows, the issue shifts. The real question becomes whether the organisation can prove what happened, show who was responsible and defend the product in law.

The key date is 9 December 2026. That is the point by which Member States must implement the revised Directive into national law. Also, it is the date from which the new rules apply to products placed on the market or put into service.

Products already in service before that date will generally continue to fall under the previous liability regime, unless later updates, upgrades or substantial modifications bring them within the scope of the revised rules.

The updated Directive expressly covers software, AI systems, updates and digital components. In certain circumstances it lowers the evidential burden for claimants, particularly where products are technically complex or relevant evidence cannot be properly disclosed.

This short test is designed to expose where that legal risk begins. Each question should be answered Yes or No only. Every No is not simply a governance weakness, it is a potential point of legal exposure.

SCOPE OF USE

Y N Does your organisation use software, AI systems or third-party models to produce outputs, decisions, recommendations or actions that affect natural persons?

This establishes whether the system falls into the type of use most likely to give rise to compensable harm, complaint or legal challenge.

RESPONSIBILITY FOR THE PRODUCT

Y N Does your organisation configure, fine-tune, adapt, integrate or materially modify the software or AI system rather than using it strictly as supplied?

This tests whether your organisation may move beyond passive use and into a position of product responsibility or retained control.

DECISION RECONSTRUCTION

Y N If one specific decision made six months ago were challenged tomorrow, could you identify the exact software version, model state, configuration and decision logic that produced it?

If a single challenged instance cannot be reconstructed, defending defect and causation becomes materially more difficult.

INPUT TRACEABILITY

Y N For that same challenged decision, can you produce the exact input data, prompts, parameters, rules or source records that the system received?

Without the original inputs, it may be impossible to explain why the outcome occurred.

RECORD INTEGRITY

Y N Is all evidence relating to that decision stored as one coherent and attributable record, including inputs, outputs, timestamps, system state and user or system interventions?

Where evidence is fragmented across multiple logs, systems or vendors, accountability often fails at the point of legal reconstruction.

DISCLOSURE READINESS

Y N If a court, regulator or claimant required disclosure tomorrow, could you provide the decision record, input history, output, system version and update history in a form that is clear and understandable?

The revised PLD materially increases disclosure obligations and failure to disclose can support presumptions against the defendant.

POST-DEPLOYMENT CHANGE CONTROL

Y N Can you prove exactly which updates, retraining events, patches, threshold changes or vendor modifications occurred between initial deployment and the date of the alleged harm?

The Directive expressly extends exposure to post-sale defects arising from updates, machine learning changes and cybersecurity failures.

SAFETY AND CYBERSECURITY CONTINUITY

Y N Do you have a documented process ensuring that safety-critical updates and security patches are applied to all systems currently in service?

A failure to patch known vulnerabilities may itself become part of the alleged defect.

HARM RECOGNITION

Y N Does your risk and incident framework expressly include psychological harm, loss or corruption of personal data and other non-physical damage recognised by the Directive?

The revised PLD broadens recoverable damage beyond traditional physical injury and property loss.

THIRD-PARTY DEPENDENCY CONTROL

Y N Where the system depends on external models, APIs, cloud services or related digital services, do you retain sufficient operational, contractual and evidential control to explain and defend the outcome?

Dependency without control is one of the most significant modern liability exposure points for software and AI systems.

MULTI-PARTY LIABILITY EXPOSURE

Y N If harm were caused by a decision involving your own system, a third-party model and an integrated software or cloud provider, do you know which party could be pursued for the full amount of compensation?

Under Article 12, where more than one economic operator is responsible for the same damage, any one of them may be required to pay the full claim, even if other parties were also at fault.

PLATFORM PRESENTATION RISK

Y N If your organisation hosts, embeds or provides access to third-party software or AI tools through its own interface, is it presented clearly enough that an average consumer would understand that the product is not provided by you?

If the presentation of the product, service or interface makes it appear to be your own, liability may attach as though you were the distributor or provider.

LONG-TERM EVIDENCE RETENTION

Y N Is your decision-level evidence, including system version, configuration, inputs, outputs and intervention records, retained in a form that remains retrievable and potentially usable for up to 25 years?

Certain personal injury claims may remain actionable far beyond the standard period, particularly where harm emerges slowly, which means evidence retention becomes a long-tail exposure issue.

OPEN-SOURCE COMMERCIAL EXPOSURE

Y N Where your organisation distributes, contributes to or deploys open-source software, can you show that it is not being used in a way that creates commercial product responsibility?

Open-source status does not automatically remove liability exposure if the software is integrated into commercial products, services or decision systems.

SUBSTANTIAL MODIFICATION CONTROL

Y N Can you identify every software update, retraining event, rules change or continuous-learning adjustment that materially altered the original risk profile of the system?

A substantial modification may be treated as a new product for liability purposes, which can restart the limitation period and alter responsibility.

SUPPLY CHAIN RECOURSE READINESS

Y N Where your systems depend on components supplied by smaller software providers or specialist developers, are contractual recourse and indemnity arrangements clearly defined?

This tests whether liability between economic operators has been anticipated before harm occurs, rather than left to be disputed after a claim.

READING THE RESULT

SCORE	WHAT IT MEANS
0–3 'No' answers	Your exposure may be relatively limited, provided your records and evidence are complete, dated and retrievable if challenged.
4–7 'No' answers	There are likely to be meaningful weaknesses in your ability to explain what happened, show who was responsible or defend a claim if harm is alleged.
8 or more 'No' answers	There is a serious risk that your organisation would struggle to prove what the system did, how the outcome arose and where responsibility sits if a legal challenge were brought.
Any 'No' to Questions 13 or 15	Treat as high-severity regardless of total score because these directly affect long-tail liability and limitation periods.

SUMMARY TABLE FOR THE QUESTIONNAIRE

QUESTION THEME	WHAT IT TESTS	WHY IT MATTERS	WHEN IT BECOMES CRITICAL
Defensibility (3–7)	Reconstruction of a single challenged decision	Avoids evidentiary gaps in defect and causation	Post-harm claim
Responsibility (1–2, 10–12)	Control over system configuration and dependencies	Identifies the relevant economic operator	Deployment and integration
Exposure Time (13, 15)	Long-term retention and modification logs	Manages long-tail liability and limitation periods	Long-term legal defence
Emerging Risks (9, 14)	Expanded harm definitions and OSS integration	Captures revised forms of exposure	Risk assessment and contracting

AFTER HARM: THE QUESTIONS THAT FOLLOW

When harm is alleged, the practical questions are rarely whether a policy existed. They are:

- What product was in service?
- What changed?
- What specific decision is challenged?
- What evidence survives?
- Who can be pursued?
- Can the organisation disclose the decision record?

JUSTIFICATION FOR THE 16 EXPOSURE TEST QUESTIONS UNDER DIRECTIVE (EU) 2024/2853

The 16 questions in the exposure test are designed to bridge the gap between general compliance, namely how a system should work and evidentiary defensibility, namely whether an organisation can prove what a system actually did when a specific decision is challenged in court.

Under Directive (EU) 2024/2853, the practical focus shifts away from theoretical safety controls alone and toward the ability to reconstruct individual outcomes, identify responsibility and defend the product in law.

The purpose of the questionnaire is therefore not merely to test governance maturity, but to identify points at which evidentiary control may fail after harm, complaint or legal challenge.

USE AND RESPONSIBILITY (QUESTIONS 1–2)

WHAT IT RELATES TO

Defining the product, the nature of its use and the relevant economic operator.

WHY IT MATTERS

The revised Directive expressly covers software, AI systems, updates and digital components. The historical argument that software is merely a service is no longer a reliable defence where the software produces outputs, recommendations, decisions or actions capable of causing compensable harm.

These questions test whether the organisation is using a system in a way that may give rise to liability exposure and whether it has moved beyond passive use into a position of retained control or product responsibility.

Where an organisation configures, fine-tunes, adapts, integrates or materially modifies the system, it may assume responsibilities that go beyond those of a simple end user.

WHERE IT ORIGINATES

Articles 1–4 (subject matter, scope, definitions of product, software, related services and economic operators) and Article 8 (liability of economic operators).

This is the legal basis for treating software and AI systems as products and for identifying who may bear responsibility where the system is configured, adapted, integrated, or materially modified.

WHEN IT APPLIES

From the point of deployment, putting into service, configuration or material adaptation.

DECISION RECONSTRUCTION AND EVIDENCE (QUESTIONS 3–7)

WHAT IT RELATES TO

Instance-level reconstruction, disclosure readiness and evidentiary integrity.

WHY IT MATTERS

This is the area where many organisations are most exposed. When a single decision is challenged, the issue is rarely whether policies existed in principle. The legal question becomes whether the organisation can reconstruct that specific outcome.

That requires the ability to identify:

- the exact software or model version
- system state and configuration
- decision logic or rules in force
- input data, prompts, thresholds and parameters
- timestamps and interventions
- update history between deployment and harm

Where this evidence cannot be produced or where relevant information cannot be clearly disclosed, a court may be able to infer defectiveness or causation, particularly where the product is technically complex.

The questions in this section therefore test whether the organisation can move from general assertions about system behaviour to decision-level proof.

WHERE IT ORIGINATES

Article 9 (disclosure of evidence) and Article 10 (burden of proof and rebuttable presumptions).

This is the legal basis for disclosure readiness, decision-level reconstruction and the risk of presumptions where evidence cannot be produced or the product is technically complex.

WHEN IT APPLIES

Triggered after harm or legal challenge, but the capability must exist before deployment.

LIFECYCLE AND CHANGE CONTROL (QUESTIONS 8–9, 15)

WHAT IT RELATES TO

Post-deployment control, updates, retraining, patching and substantial modification.

WHY IT MATTERS

Software and AI systems are rarely static products, they are updated, retrained, patched, reconfigured and in some cases continue to learn over time. The legal risk does not end at deployment.

A substantial modification may cause the modified software or system to be treated as a new product for liability purposes, which may affect the limitation period and responsibility analysis.

Equally, a failure to apply safety-critical updates or security patches may itself become part of the alleged defect.

These questions therefore test whether the organisation can prove how the system changed between initial deployment and the challenged outcome.

WHERE IT ORIGINATES

Article 7 (defectiveness, including safety expectations over the lifecycle of the product, cybersecurity, updates and reasonably foreseeable effects of self-learning behaviour), Article 4(18) (substantial modification) and Article 17 (limitation periods, including long-tail claims where relevant).

This is the legal basis for post-deployment updates, patching, retraining, cybersecurity continuity and substantial modifications that may alter the product's legal identity.

WHEN IT APPLIES

Throughout the full operational lifecycle of the software or AI system.

RECOGNISED HARM AND DEPENDENCIES (QUESTIONS 10–11)

WHAT IT RELATES TO

The scope of recoverable damage and dependency on related services.

WHY IT MATTERS

The revised Directive broadens the types of harm that may be compensable. Exposure is no longer limited to traditional physical injury or property damage. It may also extend to recognised psychological harm, personal data loss, corruption or other non-physical forms of damage provided for under the Directive.

At the same time, many AI systems depend on external components such as APIs, third-party models, cloud services and digital infrastructure.

Where those dependencies contribute to a harmful outcome, liability questions may extend across

multiple economic operators. These questions therefore test whether the organisation can identify and defend dependencies that materially influence the outcome.

WHERE IT ORIGINATES

Article 6 (damage, including recognised non-physical harm and data loss) and Article 4(3) (related services and digital dependencies).

This is the legal basis for including psychological harm, personal data loss, corruption, API dependencies, cloud services and external model reliance within the exposure analysis.

WHEN IT APPLIES

At design, procurement, deployment and incident response stages.

MARKET PRESENTATION AND LONG-TAIL RISK (QUESTIONS 12–13)

WHAT IT RELATES TO

Platform presentation risk and long-term evidence retention.

WHY IT MATTERS

Liability may not depend solely on technical ownership. If the presentation of a product or interface makes it appear to an average consumer that the tool is provided by the organisation, exposure may attach even where the underlying product is third-party.

This makes interface design, branding, embedding and platform presentation legally significant.

In addition, some personal injury claims may emerge only after significant delay. This creates long-tail liability exposure extending well beyond standard operational retention practices. The result is that decision-level evidence, system versions, configurations, inputs, outputs and intervention records may need to remain retrievable for very long periods.

WHERE IT ORIGINATES

Article 8(4) (presentation and platform liability, including where an average consumer may reasonably believe the product is provided by the hosting interface) and Article 17(2) (extended limitation / expiry period for later-emerging injury).

This is the legal basis for interface presentation risk and the requirement to retain evidence for potentially extended periods.

WHEN IT APPLIES

At interface design stage and throughout long-term evidence retention strategy.

SUPPLY CHAIN AND STRATEGIC DEFENCE (QUESTIONS 14, 16)

WHAT IT RELATES TO

Open-source commercial exposure and supply-chain recourse.

WHY IT MATTERS

Open-source status does not automatically remove liability exposure. Where open-source software is integrated into commercial products, services or decision systems, responsibility questions remain live.

Equally, when harm arises across a multi-party software chain, contractual recourse and indemnity arrangements become critical in determining how responsibility may be allocated between economic operators.

These questions therefore test whether the organisation has anticipated liability allocation before harm occurs rather than leaving it to post-claim dispute.

WHERE IT ORIGINATES

Article 2(2) (open-source software and the commercial activity qualification) and Article 12 (multiple economic operators, recourse and liability allocation between parties).

This is the legal basis for supply-chain liability exposure, OSS commercial integration risk, indemnities and recourse planning.

WHEN IT APPLIES

During procurement, contracting, integration and supplier governance.

THE AI ACCOUNTABILITY LIBRARY

The AI Accountability Library is an educational reference system built to explain what must be capable of being proven when an AI or automated system is later questioned.

Its purpose is not to advise organisations on what they should do, nor to provide consultancy, legal advice or implementation guidance. Instead, it explains the concepts, distinctions, evidentiary structures and legal exposure themes that become critical when harm, dispute, audit or regulatory challenge arises.

The Library is organised as a structured body of plain-English reference material covering governance, scope and limits, authority and responsibility, control and monitoring, evidence and records and legal exposure. Across these areas, it explains the conditions under which accountability can be understood, reconstructed and examined after the fact.

Particular emphasis is placed on questions that often become decisive only later: what the system was authorised to do, who had authority to intervene, what controls existed at the relevant moment and what evidence survives to explain a challenged outcome.

The material includes more than 620 evidence requirements and over 100 legal exposure triggers, cross-referenced against major instruments and frameworks including the EU AI Act, ISO 42001, the NIST AI Risk Management Framework and the revised Product Liability Directive.

It is designed for those who need to understand the meaning and implications of accountability concepts in practice, including directors, legal teams, governance professionals, insurers, auditors, investigators and regulators.

Access to the Library is available by annual subscription.

If the issues raised in this report reflect questions your organisation is beginning to face, further explanatory reference material is available. This covers governance, evidence, responsibility, legal exposure and decision reconstruction.

For further information or subscription access: russell@aiaccountabilitylibrary.com