From Principle to Practice: What Preparedness Means in Legal and Institutional Terms

By Flip Petillion

8 January 2026

The question that naturally follows from my essay “*Preparedness as an Old Societal Skill: Why AI Is Not Exceptional*” is how preparedness should be put into practice.

If preparedness is the right way to think about AI-related risk, does it mean that we are putting the cart before the horse? Should liability rules, safeguards and regulatory boundaries not be defined first?

I would argue the opposite. Preparedness is not the result of legal certainty; it is the response to its absence.

Throughout history, societies have rarely waited for perfect knowledge before organising responsibility. Legal and institutional frameworks usually develop step by step, often after concrete harm has occurred rather than in response to abstract risk. There is no reason to believe AI will be different.

1. Preparedness Begins Where Prediction Ends

Preparedness does not assume that technological development is transparent, predictable or fully controllable. It assumes the opposite.

It starts from the idea that uncertainty is structural, and that individuals and institutions must continue to function despite it.

Preparedness is therefore not about predicting every possible outcome. It is about organising access, responsibility and the ability to intervene before harm occurs.

A simple example makes this clearer. A two-year-old child who has just learned to walk will fall. One parent may want to protect the sharp corners of the coffee table to avoid injury. The other may say that falling is part of learning and that too much protection interferes with experience. Neither view is unreasonable.

The same tension exists at a societal level.

When I was a child, it was still considered normal to place a real Christmas tree in the living room with real candles attached to it. It was festive and traditional — and obviously dangerous. My mother, understandably worried — she had eight children, and I was the youngest — allowed it, but kept a constant watch.

Today, such a practice would be unthinkable in most households. Not because people have become fearful, but because experience has accumulated. Fire risks are better understood, safer alternatives exist, and tolerance for avoidable harm has decreased — a reality that is still painfully confirmed by winter incidents across Europe.

What changed was not human nature, nor the existence of risk, but our shared understanding of where protection is proportionate. Preparedness did not eliminate celebration; it reshaped it.

The same applies to AI. What seems acceptable at an early stage may need adjustment once risks become clearer and more widely shared.

Preparedness does not mean eliminating all risk, nor leaving everything untouched. It means recognising vulnerability, anticipating foreseeable harm, and making proportionate adjustments — without trying to control life itself.

AI preparedness follows the same logic. It is not about removing every sharp edge, nor about ignoring them. It is about deciding which edges require protection, for whom, and in which situations.

This is especially important for AI, where access to powerful systems is global and often depends only on payment or connectivity. The key question is no longer only how AI behaves, but who should have access to it, under what conditions, and with what safeguards.

Preparedness therefore inevitably raises questions of access — not only of liability.

2. Liability as a Trigger — Not as a Framework

It is likely that the first major legal developments on AI preparedness will arise through liability claims, especially in the United States.

Early cases already point in that direction: claims involving suicide, severe psychological harm or manipulation allegedly linked to AI use. These cases do not start from theories of AI governance. They start from concrete harm and familiar legal concepts: duty of care, foreseeability, causation and responsibility.

History shows this pattern clearly.

Tobacco liability did not emerge because society as a whole was harmed, but because individuals could demonstrate concrete damage. Importantly, this also included harm to third parties, such as children growing up in smoking households or people exposed to second-hand smoke. Liability expanded beyond the direct user to foreseeable harm to others.

In Europe, however, harm has long been understood more broadly. Public healthcare systems and social security schemes absorb the long-term consequences of smoking, alcohol use and industrial pollution. Harm is not only individual; it is also societal.

Asbestos litigation, including cases involving Eternit, illustrates this well. Responsibility only crystallised after decades, once scientific knowledge, cumulative harm and corporate awareness aligned. Liability extended beyond workers to families, communities, companies and even individual directors. Delayed and diffuse harm did not prevent accountability; it shaped it.

More recent debates on PFAS contamination, including cases involving 3M, show both continuity and change. As with asbestos, uncertainty played a role. Unlike asbestos, responses were faster — though after the damage had become visible. The lesson is not that preparedness was absent, but that it was incomplete.

AI is likely to follow similar paths. Harm will not be limited to direct users. Other persons may be affected indirectly but foreseeably. Preparedness must therefore extend beyond the individual user to the wider environment.

Comparable approaches already exist. In road traffic, responsibility is shared between drivers, manufacturers, insurers and authorities. In EU law on online intermediaries, there is no general duty to monitor, but there is a duty to act once concrete risks become known. This “no monitoring, but action upon notice” logic is likely to be highly relevant for AI.

3. Access to AI as a Governance Question

This leads to a question that is often overlooked: should access to AI be unconditional? Should powerful systems be available simply because they are paid for — or free?

Societies have never treated access to high-impact tools as neutral. Driving requires a licence. Firearms require authorisation. Certain substances are restricted or prohibited, with significant differences across jurisdictions. These limits are not based on distrust, but on the recognition that unrestricted access increases risk.

Preparedness in the AI context therefore requires thinking about access thresholds, conditions of use and context. This is not about exclusion, but about proportionality — matching access to capability, vulnerability and foreseeable impact.

Ignoring access reduces preparedness to damage control.

4. When There Is No Preparedness — and Why Luck Is Not a Model

The absence of preparedness does not always lead to disaster. Sometimes, through improvisation or sheer luck, crises are managed successfully.

The collapse of the main stage at Tomorrowland — a globally known event and an export product of Belgium — did not result in casualties, and the festival continued. The organisers were widely praised.

But governance cannot be based on lucky outcomes.

The absence of victims does not remove responsibility; it delays its assessment. Such situations almost always lead to years of litigation, insurance disputes and accountability debates.

Preparedness exists to avoid governing by miracle. It does not aim to foresee everything, but to reduce reliance on luck as a substitute for responsibility.

As the saying goes: *le parfait est l'ennemi du bien*. Preparedness is not about perfection, but about doing better than chance.

5. Preparedness Beyond Law: Legal and Technical Co-Responsibility

Recent court decisions show why preparedness must be concrete. Courts have recently been confronted with briefs in which lawyers relied on AI-generated content that cited entirely fictitious case law.

These incidents do not necessarily show malicious technology or bad faith. They show insufficient preparedness.

The problem is not the use of AI, but its use without understanding its limits, without verification, and without professional judgment intervening when needed.

This is where preparedness becomes real: knowing when a tool helps, when it misleads, and when responsibility cannot be delegated.

These cases do not call for bans or panic. They call for professional awareness, institutional guidance and enforceable standards — in short, preparedness embedded in practice.

Preparedness is not purely legal. It is often primarily technical and organisational. Legal rules rarely create preparedness; they stabilise it once it exists.

Design choices, human-in-the-loop mechanisms, logging, auditability and escalation procedures are not legal details; they are preparedness. Law intervenes when these mechanisms fail or when responsibility must be allocated.

Legal and technical responsibilities therefore meet. Lawyers and judges do not design systems, but they define how systems are assessed and used. Technical safeguards without legal accountability lack legitimacy.

Preparedness is a shared responsibility.

6. Moving Beyond Doomsday Thinking

Public debate on AI often slips into fatalism. AI is portrayed as uncontrollable, opaque and inevitably dangerous. In that narrative, preparedness is sometimes dismissed as a form of retreat — a cave to hide in while the storm approaches.

That view misunderstands what preparedness is.

Preparedness is not pessimism. It is institutional maturity. It does not deny risk, but it refuses paralysis. It accepts that uncertainty cannot be eliminated, and focuses instead on organising responsibility around it.

Societies already function this way in many domains. In healthcare, security, transport and infrastructure, institutions do not wait for perfect forecasts. They remain capable of response when things go wrong.

AI does not change this logic. It reinforces it.

Treating AI as an existential rupture encourages either fear-driven regulation. Treating AI as another high-impact activity forces institutions to do something more demanding: to apply what they already know, carefully and consistently.

Preparedness is not dramatic. It is disciplined.

7. Preparedness in a Global Regulatory Environment

Preparedness must also be understood in a global context.

AI systems are deployed across borders, often simultaneously subject to different legal regimes. Developers, deployers and users cannot meaningfully isolate themselves within one jurisdiction.

In this environment, preparedness may require readiness for the strictest applicable framework — not out of regulatory anxiety, but out of operational realism.

This raises an important question: does preparedness mean aligning practices with the highest standards, precisely because AI systems do not respect borders?

That question becomes unavoidable in light of comprehensive regulatory frameworks such as the EU AI Act. Even where such rules do not apply directly, they may influence expectations, liability standards and professional conduct elsewhere.

Preparedness, in this sense, acts as a bridge. It connects fragmented legal orders with globally deployed technologies, not by harmonising law, but by harmonising responsibility.

8. Preparedness as a Stabilising Force

Preparedness does not promise safety. It promises structure.

It does not eliminate harm. It reduces dependence on luck.

It does not deny uncertainty. It makes institutions resilient in its presence.

When AI governance is framed through preparedness rather than panic, the challenge becomes less theatrical but more serious. It requires institutions, professionals and organisations to take responsibility before harm becomes visible.

That is demanding work. It is also familiar work.

Preparedness is not a reversal of logic. It is what remains when prediction fails.

That is not a retreat. It is governance.

9. Making Preparedness Tangible: Familiar Illustrations

Preparedness remains abstract until it is grounded in familiar situations.

Meteorology offers a clear example. Institutions such as the Belgian Royal Meteorological Institute (KMI/IRM) exist to inform, warn and prepare the public. Yet even the most advanced models cannot predict everything. Local phenomena may occur unexpectedly, such as sudden icing caused by freezing temperatures combined with rainfall, as happened during the night of 30–31 December 2025 in Bruges.

These blind spots do not undermine preparedness. They define its limits. Preparedness warns where possible and accepts residual risk where prediction fails.

Public order provides another illustration. After riots during New Year's Eve 2024–2025, extensive preventive measures were taken in Brussels for the 2025–2026 celebrations. Yet preparedness itself can be affected by external factors. The announcement of industrial action by firefighters, combined with statements by union representatives emphasising the unpredictability of available manpower, raises difficult questions.

If preparedness presupposes organisational readiness, how should responsibility be assessed when essential services declare that they have no control over their own capacity? These situations show that preparedness is not only technical; it is institutional.

The same logic applies to AI.

Preparedness does not require omniscience. It requires clarity: about roles, escalation paths and intervention when things go wrong.

In the legal profession, this translates into concrete responsibilities. In earlier work, I have examined how in-house counsel, practising lawyers, judges and arbitrators can integrate preparedness into their professional use of AI tools. This includes awareness of limitations, safeguards against misuse, and procedural best practices to mitigate foreseeable consequences.

Preparedness, across all these domains, follows the same logic. It does not control the future. It organises responsibility in its presence.

In subsequent essays, I will examine the importance of AI preparedness in legal services and I will discuss how the EU AI Act contributes to preparedness.