



THE WORLD'S #1
PROFESSIONAL SERVICES PROVIDER




Trustworthy AI

Trust in AI: Built in, not bolted on

As AI becomes increasingly embedded across organisations and society, trust will be central to how confidently it can be adopted and scaled. Building trustworthiness into AI from the outset can support innovation, protect those using the technology and create competitive advantage. Achieving this requires organisations to turn the traditional approach to compliance on its head.



Artificial intelligence is developing at a pace that challenges traditional approaches to regulation and compliance. Too often, AI governance is arriving at the wrong point in the development process.¹



As a result, organisations risk wasting time, energy and money relearning something the software industry spent two decades discovering – problems identified too late are more complex and expensive to fix.

Why governance is arriving too late

Across many organisations building or deploying AI right now, governance shows up almost exclusively at the end. A model is built, trained and deployed and then goes through the ‘compliance gate’ – the bias audit, legal review and data protection impact assessment.

A flurry of compliance theatre and box ticking creates the illusion of control without embedding it into the process. Quite often, this results in the identification of multiple risks that should have been designed out earlier, the introduction of a risk register and release timelines being pushed out.

It also means frustrated development teams, delayed go-to-market strategies and costly reworks. The root cause is structural: governance is treated as a gate to pass through rather than a lens through which to build. Until this changes, organisations will continue to pay the price through delays, reworks and the erosion of trust in AI initiatives themselves.


Ahead of an International AI Summit focused on what AI can do for European competitiveness, this needs to be recognised as a sequencing problem rather than a gap in the regulatory rule book. Poorly sequenced governance can become an obstacle to the very AI innovation that regulation should enable.

¹The 2026 National Digital and AI Strategy, Digital Ireland – Connecting our People, Securing our Future

What we can learn from software development

In software development, the lesson was learned the hard way. The waterfall model followed the same order that AI governance often follows today – testing tacked onto the end and defects growing steadily more expensive to fix as more was built on top of them.

Then came 'Shift Left', which moved testing, security and quality upstream into design and development. Test-Driven Development took this further, with developers writing automated tests before writing the code itself, helping to ensure that it met expected outcomes from the outset. The principle is clear: the earlier you catch a problem, the cheaper and easier it is to fix.



Software teams also discovered that **some problems only reveal themselves at scale, under real-world conditions and with real user behaviour.**

This led to 'Shift Right', which supports testing in production through techniques such as canary releases, feature flags and chaos engineering. ²



AI governance needs the same holistic approach: safety, fairness and explainability built into design rather than bolted on later. Once deployed, systems need continuous drift detection, real-time monitoring and adversarial red teaming to identify failures that only emerge when reality meets scale.

Shift left on design. Shift right on operations. Both are non-negotiable.

2.The 2026 National Digital and AI Strategy, Digital Ireland – Connecting our People, Securing our Future

Building trust into AI by design

Building trustworthy AI requires organisations to consider how systems behave in practice and address potential problems throughout development, deployment and human decision-making.

This should involve:



1 Bias and fairness testing:

Resourced and carried out when the model is selected rather than squeezed into an existing workload after the fact.

2 Provenance and content marking:

Designed into the pipeline from the outset rather than retrofitted once a standard or regulatory obligation demands it.

3 **Safety and data protection requirements:** Written into the same specification document as functional requirements, challenged and agreed by the right people at the right time.

4 **Engineers, product teams, ethicists and legal teams:** Given a seat and real authority wherever trust decisions are made rather than being brought in to complete a checklist afterwards.

Built this way, compliance stops being an afterthought and becomes a by-product of good design and a foundation for trustworthy AI.




The commercial case for trust

Regulatory fatigue is real and understandable. Organisations building AI have watched the regulatory landscape develop for years, and every new obligation can become another line item.

If the case for catching problems early is presented simply as another argument for compliance, it risks being ignored alongside everything else competing for attention.

Set aside the compliance requirement for a moment and look at what changes when trust is built in proactively.




- 1 Speed to market:** Products designed this way can clear procurement, legal and enterprise client reviews faster because the evidence a reviewer needs already exists rather than being assembled after the fact. Due diligence is cleaner.


- 2 Customer and consumer trust:** Organisations that can demonstrate provenance, fairness testing and human oversight as designed-in features are better placed to build confidence in how they use AI.

- 3 Talent and investment:** Organisations that demonstrate a serious approach to trustworthy AI can strengthen their appeal to talent and investors. It signals that responsible AI is embedded in how the organisation operates rather than treated as a compliance exercise.

This is a commercial argument as much as a compliance one. Businesses that build trust into AI from the outset can move faster, strengthen relationships with customers and employees, and focus their best people on strategy and innovation rather than managing avoidable problems.



Compliance becomes an outcome of doing AI well rather than the goal itself.³



³The 2026 National Digital and AI Strategy, Digital Ireland – Connecting our People, Securing our Future

Ireland's opportunity

These questions carry particular weight for Ireland.

The 2026 National Digital and AI Strategy, Digital Ireland: Connecting our People, Securing our Future, sets an ambitious dual mandate for the country - become a Global Hub for Applied AI and a trusted, agile and forward-looking digital regulatory hub.



This is not a contradiction. It is a choice. ⁴

Responsible adoption, fundamental rights and societal trust can provide a foundation for sustainable competitive advantage as Ireland implements the EU AI Act.

Ireland has established the AI Office of Ireland as an independent statutory body and central coordinating authority. The country has adopted a distributed regulatory model that draws on the expertise of existing sectoral regulators, with the AI Office coordinating their activities and supporting consistent implementation. The Office will also champion AI innovation and adoption through Ireland's regulatory sandbox.

Making this work requires collaboration between regulators, systematic information sharing and access to appropriate technical and regulatory expertise.



4. The 2026 National Digital and AI Strategy, Digital Ireland – Connecting our People, Securing our Future

Making trust a competitive advantage

Ireland's position as the European base for many major technology companies gives the country significant exposure to AI development and deployment. This brings both opportunity and responsibility.

The country's competitive advantage will not come from minimising or downplaying the risks. It will come from demonstrating that rapid innovation and high levels of trust can coexist and that businesses do not have to choose between speed and safety.

Embedding trust into the development and adoption of AI is not a regulatory burden. It is a market differentiator.

For Ireland, the opportunity is to build an economy capable of attracting and retaining AI talent, investment and capability precisely because businesses and users have confidence in the regulatory environment and the principles that govern it.



Contact



Nicola Flannery

Partner

Technology & Transformation
(Trustworthy AI)

niflannery@deloitte.ie

Deloitte.

At Deloitte, we make an impact that matters for our clients, our people, our profession, and in the wider society by delivering the solutions and insights they need to address their most complex business challenges. As the largest global professional services and consulting network, with over 470,000 professionals in more than 150 countries, we bring world-class capabilities and high-quality services to our clients.

In Ireland, Deloitte has over 3,000 people providing audit, tax, consulting, financial advisory, and risk advisory services to public and private clients spanning multiple industries. Our people have the leadership capabilities, experience and insight to collaborate with clients taking them wherever they want to go.

This document has been prepared by Deloitte Ireland LLP for the sole purpose of enabling the parties to whom it is addressed to evaluate the capabilities of Deloitte Ireland LLP to supply the proposed services.

This document is not an offer and is not intended to be contractually binding. Should this proposal be acceptable to you, and following the conclusion of our internal acceptance procedures, we would be pleased to discuss terms and conditions with you prior to our appointment and no reliance may be placed for any purposes whatsoever on the contents of this document.

Deloitte Ireland LLP is a limited liability partnership registered in Northern Ireland with registered number NC001499 and its registered office at The Ewart, 3 Bedford Square, Belfast, BT2 7EP, Northern Ireland.

Deloitte Ireland LLP is an affiliate of Deloitte EMEA BV ("EMEA"), which is a member firm of Deloitte Touche Tohmatsu Limited ("DTTL"). EMEA and DTTL do not provide services to clients. Services may be provided by the EMEA shareholders or their affiliates, which are separate and independent legal entities. Please see www.deloitte.com/about to learn more about our global network of member firms.

The Deloitte Ireland LLP privacy notice is available at www.deloitte.ie/privacy

© 2026 Deloitte Ireland LLP. All rights reserved.