

AI Governance Webinar Series: Workplace, Risk Assessments & Legal

June 2026

Table of Contents

Introduction.....	4
1. Governing AI in the Workplace in Practice	5
<i>The Current State: How AI Is Already Embedded in Your Workplace</i>	<i>5</i>
<i>The Most Common Risks and How to Address Them</i>	<i>6</i>
Over-reliance and Habituation.....	6
Cultural and Language Bias.....	6
Invisible Automation: The Silent Red Flag	6
<i>The Regulatory Landscape: What You Need to Know.....</i>	<i>6</i>
<i>Governance in Practice: Concrete Approaches.....</i>	<i>7</i>
Build a Use Case Registry	7
Risk-Tiered Monitoring.....	7
Training and Awareness That Actually Works.....	7
Reframe What Governance Is For	7
<i>Closing Advice: The Panel's Key Takeaways</i>	<i>7</i>
2. AI Risk Assessments in Practice	9
<i>What Is an AI Risk Assessment?</i>	<i>9</i>
The One-Stop Shop Decision	10
<i>Scaling Risk Assessments Across the Organisation</i>	<i>10</i>
Collect Less, But Better	10
A Two-Tier Approach.....	10
Embed Into Existing Processes.....	10
Using AI to Govern AI.....	11
<i>Avoiding Duplication: One-Stop Shop vs. Separate Assessments</i>	<i>11</i>
<i>Living Assessments: Triggers for Re-evaluation.....</i>	<i>11</i>
<i>The EU AI Act and Regulatory Preparedness</i>	<i>11</i>
<i>Closing Advice: The Panel's Key Takeaways</i>	<i>12</i>
3. AI in Legal: Contracts, Liability, and Risk.....	13
<i>Setting the Scene: Why Legal Sits at the Heart of AI</i>	<i>13</i>
<i>From SaaS to AI: Why the Old Contracts No Longer Fit.....</i>	<i>14</i>
Data use and hidden terms.....	14

IP and ownership of outputs	14
Performance and accuracy.....	14
Use case is everything.....	14
Borrow the language of technical standards	15
<i>Liability: A Patchwork Under Pressure</i>	15
<i>Public-Facing Chatbots and the Limits of Disclaimers</i>	15
<i>Emerging Issues: Building Skills, MCPs and the EU AI Act</i>	16
<i>Closing Advice: The Panel's Key Takeaways</i>	16
4. How White Label Consultancy and Pritect Can Help	18

Introduction

Artificial intelligence is already embedded in the tools employees use every day, the systems that shape decisions about people, and the contracts that govern how companies buy and supply technology. For most organisations, the question is how to govern it well enough to remain accountable when something goes wrong.

Over three webinar sessions, White Label Consultancy and Pritect convened practitioners, lawyers and governance professionals to explore that challenge from the inside. Each session focused on a distinct dimension of the problem: how AI is being governed in the workplace today; how organisations are scaling formal risk assessments as deployments multiply; and what the legal and contractual landscape looks like as the technology outpaces the frameworks built to contain it.

This white paper captures the key insights, practical frameworks and panel recommendations from all three sessions. It is organised by session, so readers can go directly to the topics most relevant to them, or read end to end for a comprehensive view of the current state of AI governance in practice.

The conversations were candid about the scale of the challenge. Regulation is a global patchwork that is still taking shape. The contracts and processes organisations relied on even a few years ago no longer fit the technology they are now deploying. And most governance programmes are under pressure to do more with less. What the panelists consistently offered, however, was not alarm but actionable direction: approaches that work in practice, common mistakes worth avoiding, and a clear sense of where to begin.

1. Governing AI in the Workplace in Practice

Speakers

- **Ra'ad Siraj:** Trusted Data and AI Executive / Global Ecosystem Connector
- **Sarah Mathews:** Responsible AI Manager / Adecco
- **Roy Kamp:** Legal Director / UKG

Moderator

- **Federico Marengo:** Associate Partner / White Label Consultancy



The Current State: How AI Is Already Embedded in Your Workplace

One of the initial themes from our webinar was this: most organisations significantly underestimate how deeply AI is already embedded in their day-to-day operations. It goes well beyond the chatbots we hear about in the news. AI is now active in scheduling, productivity tracking, and performance management; safety monitoring and customer service routing; promotion and compensation decision-support; and everyday communication such as drafting emails and translating documents.

Much of this adoption is employee-driven, not employer-mandated. Staff independently use general AI tools, and these grassroots deployments carry their own risks that organisations need to account for. The real inflection point comes when AI moves from a pilot to simply how work gets done. At

that moment, accountability becomes diffuse. Nobody remembers who approved the system or what to do when things go wrong. This is why governance needs to be in place before that transition, not after.

The Most Common Risks and How to Address Them

Over-reliance and Habituation

After a while, people stop checking AI outputs because the system has always seemed right before. This is habituation, not trust, and it leads to consequential decisions being made without genuine human review. Ra'ad's recommended fix is to introduce friction by design: build deliberate checkpoints into high-stakes processes such as hiring, discipline and performance that require a human to actively confirm a decision, not just accept it by default. This maintains oversight without slowing down operations.

Cultural and Language Bias

Most large language models are trained predominantly on Western, English-language content. Sarah highlighted how this quietly shapes how people write and communicate globally, creating subtle but real cultural misalignment across diverse workforces.

Invisible Automation: The Silent Red Flag

Ra'ad identified key early warning signs: employees who stop asking why a decision was made, managers who say 'the system flagged it, so we acted', and governance teams that stop receiving questions because everyone assumes the AI is fine. When these signs appear, your governance has become invisible, and that is when risk escalates most rapidly.

The Regulatory Landscape: What You Need to Know

Roy outlined a critical reality: the regulatory landscape is a global patchwork. Organisations operating across multiple countries must map the frameworks relevant to them. The key ones raised in the session were GDPR, the general baseline in the EU covering automated decision-making with significant effects on individuals; the EU AI Act, with more specific rules for high-risk AI use cases whose core principles remain stable despite ongoing reviews; US state laws, where California, New York, Illinois and others have enacted specific AI regulations with no single federal standard yet; and employment and labour law, which varies significantly by country and intersects directly with AI use in hiring, performance and discipline. The overarching principle is that the higher the stakes of a decision, the more robust the human oversight needs to be.

A practical mindset shift from Roy: regulators are not asking organisations to slow down. They are asking organisations to show their working, to document decisions, reasoning and safeguards consistently. Roy recommended proactively contacting your local regulator if you are uncertain about a high-risk use case, as many data protection authorities offer anonymous helplines for this purpose. Sarah added an ethical baseline: understanding who your system affects, testing it properly, and being willing to act on what you find matters more than ticking boxes.

Governance in Practice: Concrete Approaches

Build a Use Case Registry

Policies alone do not work. The solution is a registry that every AI deployment must pass through before going into production. The act of registration forces a structured conversation about what the system does, who it affects, what data it uses, and what level of oversight is required. Critically, Ra'ad stressed, this intake process should be embedded into existing workflows such as software development lifecycles, vendor procurement and partner agreements, so governance happens naturally rather than as a separate bureaucratic step.

Risk-Tiered Monitoring

Not everything needs the same level of scrutiny. High-stakes systems such as hiring, discipline, compensation and safety warrant continuous monitoring, regular audits and mandatory human review, while lower-stakes systems can rely on periodic spot checks and anomaly detection, with a focus on user enablement. A crucial insight from Ra'ad: all AI systems degrade over time as the world changes and training data becomes stale. Monitoring is what tells you when a system needs to be replaced, not just patched.

Training and Awareness That Actually Works

The panelists shared what actually works in practice: role-based training, so people know what AI means for their specific job rather than AI in the abstract; short and frequent sessions, because the technology moves too fast for long, static modules; and ambassador programmes, where peer educators spread good practice and surface real feedback.

Reframe What Governance Is For

This one is perhaps the most important. Ra'ad explained that organisations should embed oversight into product and engineering teams from the start, not bolt it on at the end. People work around governance when they think it exists to protect the company from them. When they understand it exists to protect them, so they are not the person standing alone in front of a regulator when something goes wrong, the dynamic shifts from adversarial to collaborative.

Closing Advice: The Panel's Key Takeaways

- **Treat governance as a product, not a process.** A process is something you follow; a product is something people want to use because it makes their job easier (Ra'ad)
- **Run pilots and actually listen to feedback.** You will learn more in a few weeks of observation than in months of policy writing (Sarah)
- **Do not be afraid of compliance.** Partner with it. Legal and ethics teams exist to get you to yes, not to say no (Roy)
- **Apply AI to your governance programme itself.** If you are not using AI to help scale your oversight work, you will not keep pace with the rate of deployment (Roy)
- **Know who your system affects before you deploy.** That is the foundation of meaningful testing and real accountability (Sarah)

If the themes from this session have prompted questions about how to put governance into practice at scale, Section 4 explores how White Label Consultancy and Pritect can help.



2. AI Risk Assessments in Practice

Speakers

- **Cristina Ulessi:** Head of AI Governance / Bupa
- **Benjamin Laroche:** Senior Legal Counsel / Airbus
- **Noemie Weinbaum:** Senior Managing Counsel / UKG

Moderator

- **Federico Marengo:** Associate Partner / White Label Consultancy



What Is an AI Risk Assessment?

One of the first themes from our webinar was the need to distinguish AI risk assessments from other familiar instruments, such as DPIAs. Cristina opened with a clear framing: a good AI risk assessment seeks to do two things, first identify and then mitigate the risks linked to responsible AI principles such as fairness, safety, robustness, transparency and explainability. She outlined two main dimensions the assessment must cover. The data and model dimension looks at what data is used for training, validation and testing, what the model is and how it performs, and who the intended users are. The practices dimension looks at how users are guided to use the system within its intended scope, how over-reliance is avoided, and how human oversight and performance monitoring are ensured.

Cristina also noted that assessments look quite different depending on whether the solution is developed in-house or procured from a third party. In-house development gives direct access to technical detail, while procurement requires reliance on vendor disclosures, which may not reach the level of depth you need.

The One-Stop Shop Decision

A key design question for organisations is whether to combine the AI risk assessment with other assessments, such as DPIAs and information security risk assessments, into a single process, or to keep them separate. Cristina's view was that there is no universally right answer; it depends on where AI governance sits in the organisation. But regardless of the approach, the assessments must be coordinated and coherent. An AI risk assessment that contradicts the legal basis decisions made in a DPIA creates more problems than it solves. She also emphasised that AI risk assessments are living documents: use cases change, scope evolves, and a pilot eventually becomes a production system, so assessments must be revisited as circumstances change and recommendations updated accordingly.

Scaling Risk Assessments Across the Organisation

The panelists spent considerable time on what is often the most pressing practical challenge: how do you manage not one assessment, but tens, hundreds, or thousands of them?

Collect Less, But Better

Benjamin shared that the AI projects requiring assessments at large organizations can run into the hundreds. His core insight was to resist the temptation to collect every single detail. Organisations often build intake processes with too many questions in the hope of generating useful data later; in practice, the result is poor quality responses, innovation bottlenecks, and analysis that never happens. His recommendation was that five to six targeted questions are sufficient for the first triage, enough to determine risk level and prioritise further work, and he suggests to be strict on requests to add questions unless there is a clear, immediate use for the data.

A Two-Tier Approach

Benjamin described a two-tier model that large organizations can find workable. Tier one is a short self-screening triage that determines risk classification, whether prohibited, high risk, limited risk or low risk, using a deterministic formula. Tier two is a more detailed assessment for cases that need it, distributed to the relevant legal, governance and technical teams. This allows resources to be directed where they matter most and ensures compliance does not become a uniform burden applied to everything equally.

Embed Into Existing Processes

Both Benjamin and Cristina emphasised the importance of not adding yet another compliance layer on top of existing ones. Benjamin noted that in many large organizations new AI or IT projects already may pass through ten or more compliance checkpoints covering privacy, data governance, export control and, depending on the sector, national security. The practical approach is to embed AI governance into those existing workflows rather than treating it as a separate bureaucratic step.

Using AI to Govern AI

Noemie suggested to use AI-powered “data detective agents” to help scale the review process. These agents can be fed with internal policies, processes and applicable regulations ingest tickets raised by engineers and third-party risk teams and flag issues for legal review. This would mean triaging at scale so the legal team can focus attention where it matters most. Her broader point: it is not possible to govern AI at scale without using AI itself, as manual review of every intake is simply not feasible.

Avoiding Duplication: One-Stop Shop vs. Separate Assessments

The group returned to the question of coordination between different assessment types, with a focus on avoiding duplication and assessment fatigue. Cristina’s practical suggestion was to establish a cross-functional committee with representation from privacy, information security, legal and AI governance. This creates visibility of what other functions are assessing, ensures recommendations are coherent, and avoids asking the business the same questions multiple times.

Benjamin added that AI governance teams may also be dealing with employees building their own AI tools, such as custom agents and templates shared across teams, and emphasised that these need governance too. A potential approach is to establish a light-touch approval step at business-unit level before anything goes through the formal compliance process. On the supplier side, he noted that large organizations may consider building compliance requirements directly into vendor contracts, including the ability to monitor AI systems for accuracy drift over time, either through the vendor’s own tools or via an API.

Living Assessments: Triggers for Re-evaluation

The panel agreed that the hardest part is not doing the first assessment, it is knowing when to do it again. A review should be triggered when the system is being used in a new way or rolled out to a wider group, when it is moving from a pilot into full production, or when an incident or performance issue flags something that needs a closer look.

Cristina’s view was that this only works if your AI inventory is kept live and up to date, and if people have a real reason to maintain it rather than just a compliance obligation. If the inventory reflects what is actually happening, it becomes much easier to spot when a reassessment is needed.

The EU AI Act and Regulatory Preparedness

The panellists were asked how organisations are approaching the EU AI Act, especially given the ongoing uncertainty around timelines in April 2026. Noemie’s observation was that companies are not waiting for the regulation to be finalised before acting; their enterprise customers are already asking for transparency and contractual commitments around AI governance, and the market is moving faster than the regulation.

Benjamin was direct about what this means in practice. He mentioned that AI governance professionals and legal teams regularly encounter suppliers that arrive at contract negotiations with AI certifications dated 2023, which no longer reflect the current state of the technology or governance expectations. What is needed is evidence that a supplier can help you meet your obligations today, not a document that was accurate three years ago. His practical warning: some

obligations, like post-market monitoring for high-risk AI, take around two years to properly implement at large organisations, so if you have not started planning, you are already behind.

Closing Advice: The Panel's Key Takeaways

- **The topic is already complex**; over-engineering your process will make it unmanageable. Start simple and build from there (Benjamin).
- **Build processes and tools that can follow the movement of the technology** rather than ones that assume a fixed target (Noemie).
- **Design your assessments as if you will need them again**, because you almost certainly will. That mindset is the foundation for your scaling infrastructure and guardrails (Cristina).

The challenges this session surfaced (managing hundreds of assessments across overlapping systems, keeping them live, and scaling without drowning in manual effort) are exactly the problems that structured tooling and specialist advisory are built to solve. For a closer look at how White Label Consultancy and Pritect approach this, see Section 4.



3. AI in Legal: Contracts, Liability, and Risk

Speakers

- **Peter Hense:** Lawyer / Spirit Legal
- **Yelena Ambartsumian:** Founder / Ambart Law
- **Olivia Ward:** Senior Associate / Burges Salmon

Moderator

- **Federico Marengo:** Associate Partner / White Label Consultancy



Setting the Scene: Why Legal Sits at the Heart of AI

This was the third webinar in White Label Consultancy's series, following earlier sessions on AI in HR and AI risk assessment. Legal is unavoidably one of the central issues in AI, because almost every engagement like procuring a tool, supplying one, or building your own runs through a contract. The session set out to unpack what belongs in those contracts, how liability should be shared, and what both companies and law firms need to consider before they deploy. A clear thread ran throughout: the documents and habits inherited from traditional software no longer fit the technology we are now buying, and almost every practical problem the panel raised traced back to that mismatch.

From SaaS to AI: Why the Old Contracts No Longer Fit

Yelena walked through the shift. Traditional SaaS contracts assumed deterministic software: the same input gave the same output, so a service level agreement could simply guarantee uptime. Generative AI broke that assumption. As Yelena explained, these systems are probabilistic, the same input can produce many different outputs, which may even be shared with other users, so the promises that matter shift away from availability and toward accuracy. A system can be available all the time and still be wrong for your purpose. She traced the progression further: from generative AI to AI agents, where the software acts and the human supervises, and on to agentic AI, which can set and pursue its own goals, moving from a human "in the loop" to a human "over the loop." Each step forces the indemnification, liability and audit provisions to change again.

Olivia Ward sharpened the practical point. Teams reach for the software contract they already have and try to map the AI use case onto it, but a SaaS template assumes the software does what it says, you spec it, pay for it, and it works. Generative AI does not behave that way: the output varies, and the vendor often will not warrant it.

Peter Hense added the European dimension. The major shift in the EU is conceptual: AI systems are moving from being treated as software or a service to being treated as a product. As of the end of this year they fall under the product liability regime, which rests on a no-harm principle and makes you responsible across the whole lifecycle. As Peter put it, whether you develop and sell or deploy, your job is never over, obligations run in many directions, and a judge looking back will ask whether you did everything you could to prevent the harm.

Data use and hidden terms

Olivia singled out data use as the clause she spends the most time on and the one templates most often get wrong. The vendor's rights to use your inputs are frequently not in the main agreement at all, but tucked into an acceptable use policy, supplemental AI terms, or a privacy policy incorporated by reference. Buried there, vendors quietly take broad rights to use customer inputs to train or improve their models. Clients need to hunt down those supplemental terms and understand what they have actually agreed to.

IP and ownership of outputs

Olivia's second concern was ownership of what the AI produces. If a drafting tool's output ends up in a contract and the model was trained partly on third-party content without proper licences, the chain of title becomes murky. Most vendor terms disclaim any warranty on outputs and say nothing about ownership, which is not good enough when the output is commercially significant or will be published or relied upon.

Performance and accuracy

Her third point was performance. Software contracts have SLAs for uptime, response time and bug fixes; AI contracts often have no equivalent for output quality, and no contractual definition of what "good" looks like. That makes it hard to argue breach when a model produces poor outputs or degrades over time. Olivia advised pushing hard for an accuracy obligation or at least a fitness-for-purpose standard.

Use case is everything

Yelena stressed that the sticking points depend heavily on the use case. Many US AI laws in force today are effectively disclaimer laws, while older legacy laws some decades old, with robust case

law still apply. Where a vendor agreement pushes all compliance obligations onto the client, especially in sensitive areas like consent recording or consequential decision-making, those are some of the provisions legal teams should push back on and try to carve out.

Borrow the language of technical standards

Peter offered a practical tip for moments when a lawyer is unsure about the technical aspects of AI systems in contracts: lean on technical standards and terminology from bodies such as ISO, IEC and IEEE. He described asking providers to supply an AI system impact assessment; if they can produce one, they have the documentation ready and have already tested the system, and it can be folded into the buyer's own quality management process.

Liability: A Patchwork Under Pressure

Liability sits in every contract, it decides who pays when things go wrong, and the panel agreed the traditional approach is straining. Yelena described the trajectory: early on, when ChatGPT first appeared, caps were borrowed straight from SaaS contracts, capped at the fees paid in the previous twelve months. With AI agents and agentic platforms, buyers' lawyers have grown more sophisticated and pushed back. The market is moving toward a multiple of annual fees, with carve-outs and much higher caps or no cap at all for gross negligence, wilful misconduct, data security incidents and IP infringement, with IP infringement increasingly treated outside the standard cap.

Olivia explained why the old architecture does not fit. Software failures tend to be discrete and detectable, whereas AI failures are more systemic and slow, and may not be noticed immediately. By the time anyone sees the problem, harm can have accumulated in a way a twelve-month fee cap never anticipated.

Peter stated that several competing legal systems collide here, and in Europe you are sharply limited in how far you can exclude liability. One increasingly practical answer, he suggested, is insurance, something you can actually negotiate now, since no one can be sure how a court will read a given clause in five years. Insurers impose their own conditions, which loops back to whether the buyer has real risk management and governance, including a plan for decommissioning systems at end of life.

Public-Facing Chatbots and the Limits of Disclaimers

Federico raised the familiar problem of customer-facing chatbots that hallucinate and make representations on a company's behalf, citing cases like Air Canada and DPD. The panel's advice converged on design rather than paperwork.

Yelena described building hard stops into products. A chatbot acting as a receptionist for a healthcare provider, for example, must decline to give medical or legal advice and stay limited to scheduling. Much of the protective work, she noted, happens on the product side, architecting the system to avoid trouble rather than relying on a disclaimer that simply says "AI can make mistakes."

Olivia recommended running a risk assessment before negotiations even begin: ask what the worst realistic outcome is if the system fails, how quickly you could detect it, what remediation would cost, and whether the liability cap actually covers that.

Peter warned that in the EU a disclaimer often does little. As he put it, telling someone the chatbot might be wrong is like warning that you might hit them and then doing it, you remain liable either way. He urged precision about what the system is, where it ends, and what its intended purpose is, because most products fail on unrealistic expectations.

Emerging Issues: Building Skills, MCPs and the EU AI Act

Asked how lawyers can get up to speed, the panel was consistent. Olivia put curiosity first, read newsletters, play with the tools, build a network, and attend sessions like this one. Yelena, drawing on her litigation background, compared it to preparing to depose an expert: you learn from the people closest to the work, often the data engineers rather than the CTO. Peter pointed to technical standards as the most reliable starting point in a field crowded with hype, noting that marketing language shifts while technical language stays, giving cross-border teams a common reference.

On the newest frontier, Yelena described working on an AI addendum for an MCP (Model Context Protocol) integration. The fresh concerns include credential breach notification, not just data breaches, tool definition change management, agent classification, and liability that now spans the customer, the MCP provider and the model provider.

On the EU AI Act's likely postponement, Olivia observed that many companies are taking a wait-and-see approach, much as they did before GDPR, using the extra runway to defer hard conversations that will inevitably return because a vendor's compliance position does not automatically protect the buyer. Peter was emphatic that the Act is not "GDPR 2.0" but product safety law, requiring a conformity assessment and auditable documentation before a system reaches the market; without it, a competitor could seek an injunction and shut a system out of the European market. Beneath the AI Act, he argued, the real long-term influence is the growing landscape of harmonised standards developed with ISO and IEC, and given product cycles, compliance work should start now.

Closing Advice: The Panel's Key Takeaways

- **Read the data use provisions closely.** Find where they live, the MSA, referenced documents, or a live link that changes through the year, and ask every question, even ones that feel basic, to understand exactly what you are buying (Olivia)
- **Don't assume anyone knows everything.** Things move so fast that no one is truly an expert; if a contract term is unclear, ask your counterparty's counsel to walk you through it rather than accepting sales-team answers (Yelena)
- **Don't buy AI just to buy AI.** Define your actual problem first and ask whether you genuinely need AI to solve it, that alone avoids most of the reasons AI projects fail (Peter)
- **Architect for safety, don't rely on disclaimers.** Build hard stops and guardrails into the product, and run a risk assessment before you negotiate or deploy (Yelena and Olivia)

If the themes from this session have prompted questions about how to put governance into practice at scale, the next section explores how White Label Consultancy and Pritect can help.



4. How White Label Consultancy and Pritect Can Help

Across all three sessions, the same practical challenge surfaced: AI governance is a question of operating a programme. Systems are already embedded in day-to-day work, risk assessments multiply into the hundreds, contracts and liability models are shifting under your feet, and regulations like the EU AI Act are arriving faster than most organisations are ready for. Doing this well, and being able to prove you have done it, is difficult to sustain on spreadsheets, email threads, and good intentions.

This is where [White Label Consultancy](#) and [Pritect](#) work together: expert advisory to set the right direction, and an operational platform to run it at scale.



White Label Consultancy brings the advisory and regulatory expertise. Our team supports organisations across data protection, AI governance, and cybersecurity, helping you interpret obligations under the EU AI Act, GDPR, and the wider global patchwork, design governance frameworks that fit how your business actually works, and translate legal and ethical requirements into decisions your product, engineering, and procurement teams can act on. Whether you are standing up your first AI governance function or maturing an established one, we help you get to a defensible position rather than a box-ticking one.



Pritect is our AI governance, data protection, cybersecurity and compliance platform, built to turn that strategy into a living, scalable operation. It is built by practitioners, for practitioners, and designed to address exactly the problems and the panelists described:

- **Use-case registries and intake:** capture every AI deployment through a structured, lightweight intake that forces the right conversations early, embedded into your existing workflows rather than bolted on as another bureaucratic step.
- **Risk-tiering and triage:** classify systems by risk level so high-stakes deployments get continuous oversight while lower-risk ones get proportionate, efficient review.
- **Living assessments:** keep assessments current as use cases evolve, pilots move into production, and incidents surface, with a clear view of when re-evaluation is triggered.
- **Evidence and audit-readiness:** maintain the documentation, decisions, and safeguards that the EU AI Act's conformity expectations and product liability regimes increasingly demand, ready to show your working when it matters.
- **AI-assisted scale:** use AI to help identify, evaluate, and mitigate risks across providers and deployers, so your team can focus attention where it matters instead of drowning in manual review.

The panelists were consistent on one point: you cannot govern AI at scale without the right processes and tooling, and the organisations that start building that infrastructure now will be the ones still standing when the regulation and the scrutiny, fully arrive.

If the themes from this series resonated with you, whether you are putting your first registry in place or managing your hundredth assessment, we would be glad to explore how White Label Consultancy and Pritect can support your organisation.

Get in touch: hello@whitelabelconsultancy.com ·

www.whitelabelconsultancy.com

www.pritect.ai