The Global Compliance Hub

One jurisdiction is not a compliance strategy

The EU AI Act is not the last AI regulation that will matter to your organization. It is the first one that carries material financial penalties. The others are already live.

GDPR has been in force since 2018. Japan amended its personal information law in 2022 and added an AI Promotion Act in 2025. California's privacy regime has been enforced for years. Australia's Privacy Act carries penalties up to fifty million Australian dollars for serious interference. The UK GDPR, interpreted and enforced independently since Brexit, has produced significant actions against organizations that assumed it was identical to its EU predecessor and were wrong. Any organization deploying AI at commercial scale in 2026 operates under several simultaneous regulatory obligations. The question is not whether those obligations exist. It is whether the organization can see them.

The multi-jurisdiction reality

Consider a company that is not unusual: an EU-headquartered financial services firm with customers in Japan and the UK and institutional clients in California, running an AI system that assesses creditworthiness. That single deployment sits under the EU AI Act as a high-risk system, under GDPR for the personal data it processes and its automated-decision disclosure duties, under Japan's regime for its Japanese customers, under the UK GDPR as interpreted independently of the EU, and under California and US federal guidance for its American clients.

These frameworks do not agree. One requires a continuous risk management process throughout the system's lifecycle. Another requires a structured, point-in-time impact assessment before processing begins. A third emphasizes transparency and human oversight using definitions that do not map cleanly onto the others. The organization's compliance problem is not one problem. It is five overlapping problems, each with a different authority, timeline, penalty structure, and evidentiary expectation.

Qaori.Tech OÜ, Sepapaya 6, 15551 Tallinn, Estonia, Register 17169588.
Qaori UK LIMITED, 66 Paul Street, London EC2A 4NE, UK, Register 16952497.
Qaori AI LLC, 1207 Delaware Ave #1634, Wilmington, DE 19806, United States, Register 61-2310301
Qaori AI Japan, Minato Mirai, Yokohama 220-0012, Japan. (Launch Q3 2026)
Qaori AI Inc., 1207 Delaware Ave #1634, Wilmington, DE 19806, United States, (Registration ongoing)
Qaori AI Tecnologia Brasil Ltda., Rio de Janeiro. (Launch Q4 2026)
Qaori.Coop is a service operated by Qaori Tech OÜ, Estonia, registered as Qaori DAO LLC registered at 108 West 13th Street, Wilmington, New Castle, Delaware, 19801, USA.

© 2026 Qaori.Tech OÜ. All rights reserved. Patents Pending - USPTO Provisional Application No. 64/084,387

This is why a tool that solves one of them does not materially reduce the risk. It reduces the visible risk, which is more dangerous than visible non-compliance. Non-compliance across the board is at least a coherent problem with a coherent fix. Partial compliance creates the confidence that one has governed while leaving the most consequential gaps invisible. The gap that matters is not between compliant and non-compliant. It is between visible and invisible risk. Visible risk can be managed. Invisible risk accumulates.

Why the existing tools fail

Most AI compliance tools were built for one framework, usually the EU AI Act, because it has been the most urgent. That was a rational choice when those tools were built, and it is an increasingly inadequate one now. Their structural weakness is not that they get their home jurisdiction wrong. They have no model for what happens when their home jurisdiction's obligations interact with obligations from outside it. They do not surface the conflict. They will not tell a client that satisfying one framework's logging-retention requirement can collide with another framework's data-minimization principle. They produce a document stating compliance with one regulation and leave the client to discover the cross-regulatory gap through enforcement. A tool that can only answer questions inside its declared jurisdiction is not a compliance counsel. It is a checklist, and a checklist that omits the question is more dangerous than no checklist, because it provides the confidence of a completed assessment while leaving the real gaps unaddressed.

Static question-tree wizards fail for a related reason. A fixed tree handles the expected case and routes everything else to a gap it cannot address: a system deployed as a component inside a third-party product, a system whose function changes after deployment in ways that change its risk class, the tension between a required human-oversight mechanism and the reality that the humans operating it may not understand it well enough to exercise it. These are not obscure edge cases. They are the normal operating conditions of real AI deployments. The value of a counsel is precisely the capacity to reason about the question that was not anticipated, to surface the obligation the client did not think to look for, and to ask the question the client did not know to ask. That cannot be encoded in a question tree.

Document versus evidence

A compliance assessment that produces a document proves that someone assessed something at a point in time. A system that produces evidence proves that compliant behavior is occurring. For the frameworks that govern enterprise AI in 2026, only the second matters.

The obligation is increasingly explicit that risk management must be continuous across the system's lifecycle, not a periodic review. A document cannot satisfy a continuous obligation, because a document is a record of a completed act that begins aging the moment it is produced, while the system it describes keeps changing through model updates, new integrations, and shifting context. Producing the document is the easy part. Building the infrastructure to produce continuous, verifiable evidence of compliant behavior is the hard part, and it is the part that document-generation tools have no answer to. Qaori's answer to the evidence half of the problem is described in the Sentinel and Preemptive Governance papers. The Global Compliance Hub is the half that determines what the evidence must prove, across every jurisdiction at once.

LEX: the multi-jurisdiction compliance counsel

At the center of the Global Compliance Hub is LEX, a compliance counsel Alter that covers six regulatory frameworks at once: the EU AI Act, GDPR, Japan's personal information and AI promotion regimes, the US California and federal guidance, the UK GDPR and Online Safety regime, and Australia's privacy and online safety regimes. It operates across them concurrently, surfacing cross-regulatory interactions and conflicts rather than treating each as an independent track.

As an Alter, LEX reasons rather than recites. It can run a structured, guided assessment for clients who want a predictable process, or a conversational assessment that adapts to the client's answers, identifies the highest-risk unknowns first, and treats ambiguity and cross-jurisdiction conflict as explicit objects of assessment rather than as failures it cannot handle. Either way it produces the same outputs: a jurisdiction-stamped report, an obligation register mapping every applicable obligation to the client's current control status, a gap-closure roadmap prioritized by deadline, penalty exposure, and remediation complexity, and structured evidence formatted to the evidentiary expectations of each jurisdiction. LEX does not generate documents that describe compliance. It produces evidence and roadmaps. Like every Alter, it cites its basis and is candid about its limits: it quantifies exposure and prioritizes remediation, and it does not predict enforcement outcomes or substitute for legal advice.

The penalty calculator makes the stakes concrete by quantifying exposure per jurisdiction against the current gap profile. The figures are large by design, because risk cannot be managed without being quantified: the EU AI Act reaches thirty-five million euros or seven percent of global annual turnover for the most serious violations, with GDPR, the UK GDPR, Japan, California, and Australia each adding their own. The calculator does not exist to provoke anxiety. It exists to give the organization a rational basis for prioritizing remediation against maximum possible exposure.

The USP: a compliance posture that stays current

Here is the key difference between the Global Compliance Hub and every static assessment: the regulatory knowledge base is updated twice a month.

Regulation is not a fixed target. Guidance is issued, enforcement positions shift, definitions are clarified, and deadlines arrive. A report produced against the guidance current in March is not the same compliance claim as one produced against the guidance current in September, and most compliance tooling hides that fact by producing a document with no version and no expiry. The Global Compliance Hub does the opposite. Every assessment is stamped with the version of the regulatory knowledge base it was assessed against, and that knowledge base is refreshed on a twice-monthly cadence across all six jurisdictions.

When a framework changes, the knowledge base version increments, every open assessment affected by the change is automatically flagged for re-evaluation, the client is notified, and the compliance posture is brought current. Compliance becomes a living record rather than a document that quietly goes stale. For an obligation that is itself continuous, this is not a convenience. It is the only honest way to claim a posture is current, and it is a USP that point-in-time assessments and single-jurisdiction tools cannot match, because keeping six jurisdictions current twice a month is an operating commitment, not a feature toggle.

Why August matters, and why it is not the point

On August 2, 2026, the EU AI Act's deployer obligations come into force, binding and enforceable with the penalty structure that makes the Act materially different from every framework before it. A point-in-time assessment conducted before that date does not satisfy a continuous monitoring obligation that takes effect on it. A document from a single-jurisdiction tool does not address the obligations from other frameworks that interact with it. What the date requires is an organization that can demonstrate, with evidence, which of its AI systems are high-risk, what its obligations are for each, its current control status against them, where its gaps are and what it is doing about them, and what its human oversight looks like in practice.

The deadline is not the point. Effective AI governance, governance that produces continuous evidence of compliant behavior across the full jurisdiction stack and keeps it current as the rules move, is the permanent operating requirement. August 2 is simply the date the cost of inadequate governance becomes financially material. The organizations ready on that date are the ones that built the right infrastructure before it.

What changes for the buyer

For a **general counsel or chief compliance officer**, the compliance surface stops being a set of disconnected single-jurisdiction assessments that each hide the gaps the others would reveal. It becomes an obligation register across every framework you operate under, with conflicts between them surfacing rather than being discovered through enforcement, and a posture that is rechecked as the rules change rather than aging in a drawer.

For **procurement and the board**, exposure is quantified per jurisdiction, remediation is prioritized against it, and the assessment carries a version stamp that says exactly what it was measured against and when, which is the difference between a defensible record and a hopeful one.

For an **investor**, the twice-monthly currency commitment across six jurisdictions is a real and defensible moat. Any vendor can produce a one-time assessment. Maintaining a living, multi-jurisdiction regulatory knowledge base and re-flagging every affected assessment when the rules move is an operating capability competitors cannot replicate by adding a feature, and it compounds as more jurisdictions and more clients come onto the platform.

Conclusion

One jurisdiction is not a compliance strategy, a document is not evidence, and an assessment that does not stay current is a liability dressed as assurance. The Global Compliance Hub answers all three. LEX reasons across six frameworks at once, surfaces the conflicts between them, produces evidence and roadmaps rather than documents, and keeps the whole posture current through a regulatory knowledge base refreshed twice a month. Compliance becomes a living state you can prove, in every jurisdiction you operate in, on the day a regulator asks.