

Introduction

What prevents you from moving faster?

Rules form the foundation of every enterprise, yet the intimate knowledge of how to apply them remains scattered across shared drives and spreadsheets, reliant upon individual expertise and memory.

In fact, according to Diligent's Global State of Legal Entity Compliance Benchmark Report — a survey of over 300 company secretaries, general counsel, and legal ops leaders across six regions — only one in five compliance teams have near real-time visibility into their own obligations, making governance decisions against an incomplete and often outdated picture of where they stand.

Without this single source of truth, your team are left carrying out manual intake and data collection, unable to distinguish between low-risk routine cases and those which require genuine human expertise until each has been thoroughly investigated.

And this process starts from the ground up, every time. Highly trained analysts remain mired in manual data collection and reconstruction, traipsing through email threads and Excel sheets for missing context — leaving little capacity for the higher-value judgment work.

The three-lines model was built on the basis that business units could own the day-to-day execution of compliance obligations, with your team in the second line, providing oversight and challenge from a distance, and internal audit forming the third line, objectively assuring that your governance, risk management and internal controls stay robust.

Yet, as regulatory complexity has exploded, that distance has collapsed, taking you, in the second line, with it.

Every new regulation or product launch triggers a manual re-check across contracts and policies; each market entry into a new jurisdiction means months of regulatory mapping, looking at licensing requirements, and coordinating with local counsel.

Multiply this across the scale of a large enterprise, with thousands of requests a month, spread across several jurisdictions and business units, with ever-mounting regulatory complexity to navigate and intermediaries to manage, and the model becomes simply unsustainable.

With each week that gets whittled away in reviews, pressure for timely answers from both wider business units and your board only intensifies. The chasm between obligations and the ability to execute has never been wider. This is the compliance paradox.

The first legal and compliance leaders to make these rules executable will be responsible for a significant transformation. Namely, the repositioning of compliance from a reactive control function seen as obstructive to business goals, to one which teams across the enterprise depend on to move faster and with greater confidence.]

[Section 1 - Why The Stakes Are Higher Than Ever]

Trying to fulfil both expectations – acting as the rigorous protective layer of an organisation, the buffer between the business and regulatory bodies, and still moving fast enough to contribute commercially – so often means that one is achieved at the expense of the other.

In their 2025 global compliance survey, PwC found that 72% of executives say rising compliance complexity has negatively impacted their company's profitability, and 76% say it's damaged their ability to establish and maintain third-party relationships [2]. Partners left waiting, deals delayed mid-review, it all traces back to one fundamental issue: legal and compliance teams are unable to be as proactive in applying the rules as they are in collecting them.

This is a disconnect well documented in KPMG's 2025 Risk and Resilience Survey, which found that while 48% of organisations now have centralised risk and resilience structures, only 26% report strong cross-functional collaboration [3]. Compliance, ironically, has almost become isolated from the rest of the enterprise.

The aforementioned PwC survey warns that senior compliance leaders are at “the risk of being irrelevant” – eventually, business units start finding workarounds, making their own decisions, compliance looped in only after the fact to rubber-stamp decisions which have already been made.

The instinct, understandably, has been to look for better tooling. Yet Regology's 2026 State of Regulatory Compliance Survey found that, while more than half of compliance teams now use AI, over 80% still rely predominantly on manual processes [5], highlighting how, despite rising adoption, the problem has remained. To understand why, we must look at what those tools were actually designed to do.

The Execution Gap

The first wave of Governance, Risk, and Compliance (GRC) software helped companies centralise their compliance knowledge, storing their internal policies, controls, risk registers, and audit trails all in one place. The next wave of compliance automation improved evidence collection for modern cloud environments, and, together, they made compliance information easier to document.

In large enterprises, however, the core limitation remains the application of these rules – interpreting the policies, routing approvals, and ensuring that every atomic action remains defensible across a multitude of jurisdictions, business departments and control frameworks.

When a request lands – “can I sign this contract, can I onboard this partner” – compliance teams still have to manually comb through dozens of overlapping policies, figure out which applied to this specific supplier, in this specific market, under this specific regulatory regime, and then perform this assessment entirely anew each time, regardless of whether the case is genuinely complex and ambiguous, or routine and low-risk.

This is where these legacy tools fall short. They are essentially digital filing cabinets, fixing the peripheral problem of having compliance knowledge all in one place.

Yet a rule is only useful if it can be applied in real-time, at the point of decision-making. That execution layer – capable of translating an enterprise's policies and regulatory requirements into consistent, auditable, defensible decisions – has remained out of reach for compliance teams, due in part to both mistargeting and the fundamental constraints of technology at the time.

The Unpredictability of AI

And so, with the emergence of AI and with the advancements in its reasoning, many compliance leaders are naturally turning toward generic LLMs and RAG tools to close that gap.

Some organisations have already moved past basic chatbot tools, whilst others are less familiar with its various forms and capabilities. Regardless of wherever you sit within that range, even those more well-versed with AI use cases may still underestimate what it's capable of end-to-end.

First, let's examine the most common types of AI being used: standard LLMs and RAG-based systems. Both, in the end, are still dependent on the same underlying mechanism to produce their output: probabilistic generation, an AI predicting the most statistically likely output, rather than following a fixed rule.

Nevertheless, Retrieval-Augmented-Generation and 'chatbot'-style tools aren't without their benefits. RAG systems can retrieve the exact clause, policy, or precedent relevant to a business request and surface it in seconds.

When used in tandem with a chatbot interface built on top, any business unit can directly ask a question regarding your policies and receive an answer immediately, without needing to go through compliance channels at all.

Though building on previous legacy tools such as GRC, neither presents true automation, still unable to pursue a case forward on their own, from an open request to a closed one, unprompted.

If you were to ask what was missing from a request, or how a specific policy might apply to it, they can often tell you within a single exchange. What they cannot do, is act on that information without a team member still doing the work to move the case along – physically retrieving missing data, tracking the case themselves and remembering the full context. Even once presented with everything required to make a decision, it still won't commit to a decision and is unable to carry it through autonomously.

More critically, and separate from its technical limitations, is the element of exposure they can introduce.

FINRA's 2026 Annual Regulatory Oversight Report named AI's tendency to hallucinate as a direct compliance risk, warning how easily a fabricated sanctions match, a nonexistent

adverse-media hit, or a misstated rationale behind a closed alert can go unnoticed and end up influencing compliance decisions [6].

But hallucination is just one such example of a deeper problem: probabilistic AI is inherently unsuitable for high-stakes compliance workflows, especially when it comes to three key tenets:

Consistency

When your compliance team screens a request for sanctions risk, or checks whether a third party meets a specific policy requirement, the answer cannot vary based on how the question was phrased or the time at which it was asked.

The problem arises when we consider how LLMs generate output, which it does by predicting statistically likely responses based on patterns in training data, of which there are quite literally billions. There is no fixed logic nor explicit, hard-coded rules which it refers to, and so two identical requests, submitted mere minutes apart, can return different answers.

Inconsistent decisions across similar cases are in themselves a control failure that regulators will flag.

Auditability

A regulatory investigation or internal audit will seek the answer to one question: why was this decision made?

Probabilistic AI can only reconstruct a plausible explanation after the fact, which in itself is just another generated output which it deems to be the most statistically plausible.

And, due to there being no defined rule that the AI checked against in the first place, there's nothing *to be* logged; no replayable decision trail that auditors can refer to and verify.

Probabilistic AI can also 'drift' in its logic over time as providers periodically retrain or update the underlying model, causing outputs to shift in ways that are hard to predict – and often go unnoticed until action has already been taken.

A threshold gets moved slightly, a definition softens, and it continues to generate output without warning. Compliance teams reliant on probabilistic AI soon find themselves in a precarious situation, passing review today and silently failing tomorrow.

Defensibility

Producing the right outcome is almost secondary; justifying the process behind your decision under concentrated scrutiny from audits and regulators is vital. And, should you fail, you need to be able to identify exactly where your process failed and why.

Regulators build a longitudinal picture of your organisation over time based on how it responds to failures when they occur, how transparent its processes are, how robust its

controls are, and whether it demonstrates a genuine understanding of its own compliance obligations.

They make a marked distinction between those who can be seen as having acted in good faith – with clear remediation plans and steps – and those who can only refer to a black box with no decision trail or explanation.

For the former, their fine or sanction will still reflect the breach, but their credibility and control framework remain intact. Those aforementioned organisations, on the other hand, face a much harder conversation, leaving regulators to conclude that there simply was no functioning compliance process at all.

The Build-It-Yourself Trap

The next logical thought may be to simply build on top of an LLM instead, with a well-built interface, refined prompts, and a person to review the answers. The problem with this approach is twofold: you still inherit the same probabilistic foundation, and, with an analyst then checking every response, you only create additional manual work.

One other such tempting workaround, one that an increasing number of enterprises are attempting, is to abandon that single interface altogether and build several specialised agents of their own. Each handles a specific 'subtask' or section of the workflow: request intake, research, drafting a response, and so on.

The issue, though, is that an individual agent only knows how to produce output for its own specific section of that process and cannot adapt to or access the full context behind it. As the business scales, and as regulations emerge, this lack of visibility soon manifests in two equally dangerous ways.

Visible Failure

Each agent is built and tested to expect input in a very specific, uniform manner, in a given format, with certain fields fulfilled, and in particular patterns.

Should, and very often when, that input deviates, even slightly – a new counterparty enters the picture, the source system gets updated, and a field is altered – the agent 'breaks', unable to process it, and holds up the entire chain. A person must then notice this failure has occurred, investigate the initial cause, manually correct it themselves, and finally re-run the process.

The overhead coordination soon spirals, too. Ensuring they pass data between each other appropriately, in the right sequence, without any data getting lost or corrupted along the way, soon becomes its own task.

With more moving parts, more handoffs, and more dependencies, these fragmented, multi-agent systems invite failure by nature.

Imperceptible Failure

A regulatory change introduces a new requirement, one that didn't yet exist when the agent's logic was being written, and now an additional check must be carried out somewhere in the process. Yet, because this check wasn't hard-coded into any existing agent's logic, they conduct the process end-to-end as usual. The case is neither surfaced nor flagged for human expert review; the job is deemed complete because, according to each agent's narrow logic, it was.

The error only surfaces later, under the scrutinous eye of the auditor or regulator reviewing your processes, who discovers the check was never actually carried out, and by which point your organisation has likely already acted on the decision: signing off a case incorrectly, onboarding a sanctioned entity, clearing an ineligible partner, and inviting avoidable exposure.

For large enterprises in highly regulated sectors, this is doubly important. On top of whatever jurisdictional variation already exists, you're also subject to deep, sector-specific regulatory bodies with their own ruleset you must abide by, creating another constant source of change.

In fact, 65% of general counsel and compliance officers stated that "changes in the regulatory environment" are a top risk, ranking it ahead of tariffs and other business concerns [8 5].

[Section 2 - Solution Intro] One Holistic Operating System for Legal and Compliance

As we've laid out, legal and compliance teams have long been made to compromise.

Whether between software that could only store and centralise documentation but was incapable of interpreting or applying it to live decisions, unpredictable LLMs that are inherently inconsistent and impossible to audit or defend, or an in-house build of siloed agents, no existing solution has been able to truly reduce workload without sacrificing control.

So what would a compliance system that made legal reasoning available in abundance, at any time, in any place, with all the relevant knowledge even look like in practice?

It starts with three distinct layers, each working together and doing a specific job – sitting above your existing GRC platform, risk registers, and policy repositories.

The Guardrails

Hand in hand with your own legal team, a dedicated team of lawyers and legal engineers takes the regulations your organisation must navigate, along with its internal policies, rules, and the specific way they're understood internally, and converts them into fixed, machine-readable logic.

This is the deterministic layer; the guardrails that determine what constitutes compliant for your specific organisation and industry: which policy applies to this specific request, what information is still missing before a decision can be made, what approvals, if any, are required, whether the case qualifies to be cleared automatically, and, if it doesn't, exactly when and to whom it needs to be escalated.

The AI Agents

Specifically built to act within these guardrails and unable to operate independently of them, these AI agents carry out the actual work.

A query can be made via whichever channel your business units already work in, whether a browser, Microsoft Teams, or email. Rather than them hunting down the right policy document or working out who to even ask, the agent then takes their business requests, checks said requests against the pre-defined logic, and gathers all of the necessary information.

If the case is low-risk and clearly covered by the rule, it's resolved immediately, ensuring that no routine decisions are delayed waiting in a queue behind higher-priority work.

However, if the rule does not cleanly resolve the case, it's immediately escalated to the right person, though not as a blank request. Instead, the agent presents the full context and, depending on the specific case, flags the relevant aspects to look at along with potential risks most worth a reviewer's attention – so your team can focus solely on the more complex and ambiguous cases without needing to assemble the data first.

And, because every decision the AI agent makes can be traced back as the direct result of checking each request against a specific, named rule, the platform records the full audit trail automatically: which rule the agent applied, what information it collected to apply that rule, the conclusion it reached and at what time.

The Enrichment Process

Oftentimes, an escalated case will be a recurring pattern that the agent has encountered, or a genuine edge case that was unanticipated when the rules were first written.

Rather than the agent deciding for itself how to handle it differently next time or stalling the entire process, the case is flagged automatically, and a legal engineer will review it and decide whether to and how the guardrails need to change.

This feedback loop is designed for your organisation to retain human oversight and governance as the only authority to control and modify the ruleset in real-time, with the detection of *where* that ruleset needs attention now automated.

At every stage, the same principle holds: law and compliance are far too dynamic and dependent on context ever to be *fully* deterministic.

The aim is never to remove or replace the compliance expert from the process, but to ensure that their expertise and time are directed only where it's genuinely needed – on the complex and ambiguous cases.

[Section 3] Transforming Compliance Into A Front Foot Function

As Paul F. Welter, Chief Legal Engineering Officer at Bayshore, describes:

“Law and compliance have always been resource-constrained....every compliance system has run on scarce human experts, and everything you see today is a compromise around it”.

Requests either pile up and business units are left waiting weeks at a time, if not months, for answers to their business requests, or they're processed closer to real time but without the correct level of scrutiny.

As the intermediaries multiply and jurisdictions jump in number, the people most likely to feel the pressure most acutely are those with the most institutional knowledge – an intimate understanding of the edge cases, regulatory nuances and what shortcuts will and won't hold up under scrutiny.

The default response, understandably, has been to simply hire more staff. Yet, as Thomson Reuters Future of Professionals Report 2026 found, 1 in 4 compliance and legal leaders were already actively weighing an exit from their current roles heading into 2026 [9], and 78% said they'd consider leaving their role for better support [10]. And, with each person who leaves, so too goes the specific institutional knowledge with them.

When we consider that the average replacement cost of a departing legal and compliance professional sits at around \$232,000 [11], and that 57.8% of compliance functions are already operating with five or fewer professionals [12], empowering your staff to stay and grow becomes all the more imperative.

The trend is clear; monotonous manual work, poor tooling and an inability to contribute meaningfully to the wider business all combine to sap job satisfaction and erode retention.

This is where making rules executable serves a far greater purpose beyond pure automation, allowing your existing team to be significantly more productive per capita and redirect their focus towards the high-impact judgement work they entered the industry for.

With the low-risk, routine cases resolved, your analysts who were kept occupied chasing documents and rebuilding context are then free to study the business, identify which markets are generating disproportionate risk, anticipate upcoming regulatory developments, and proactively spot and make explicit the policies that consistently generate confusion or friction.

For the first time, you gain full documented visibility across your own function's decision history. Because every decision traces back to a named rule and is logged the same way each time, you can ask whether a particular category of request has been handled consistently over the past year and catch any abnormalities. Vitally, before a regulator does.

Key institutional knowledge, previously siloed and slow to be actioned, becomes accessible to the entire enterprise. Yet beyond faster approvals and risk mitigation is the intelligence you're able to generate, and use to drive key commercial outcomes.

Over time, with each successful market entry, expedited acquisition and product launch, the board's perception of compliance shifts, from the cause of frequent delays to the ultimate authority on how, and how fast, the enterprise can capitalise on opportunities without exposing itself. And, for you personally, as the leader of that transformation, you become indispensable.

Conclusion

Despite mounting regulatory pressure and growing operational complexity, the underlying compliance processes remain remarkably unchanged.

And although research from EY discovered that 67% of organisations are making 'improvements' to their compliance function, they also found that only 31% are making a deliberate, strategic rethink of what compliance actually *is* within the business, questioning its relationship to the rest of the organisation over the next two years [13].

As Sally Trivino, EY's Global Forensic & Integrity Services Technology Co-Leader, puts it:

"Compliance should be a seatbelt, not a brake". [13].

What continues to hinder enterprises is the gap between deciphering what a rule requires and being able to action it, at the point of decision – and the avoidable risk and exposure which festers in those cracks.

The answer is a new infrastructure layer for legal and compliance work, where rules move from passive text to executable systems, available where the work happens, applied consistently, documented automatically, and escalated intelligently where human judgement is required. The future of compliance does not involve less law, but better law – in operation.