

2026

AI in Risk & Compliance

Contents

Executive summary	3
The current state of AI in risk and compliance	4
Commercials	5
Deployment	9
Risk	16
Building the AI stack in risk and compliance	25
Financial Crime Compliance	26
Data Risk	29
Cybersecurity	34
Technology Risk	37
Compliance Management	44
Conduct Risk	55
ESG	62
Market direction: The future state of AI in risk and compliance	65
Market map: AI in RegTech	69
Appendix	70
Appendix 1: Survey scope (financial services)	71
Appendix 2: Survey scope (vendors)	72
Appendix 3: Regulatory Prescription and Maturity Methodology	73
Appendix 4: Regulatory Domains Taxonomy	75
About this research	79



Executive summary

AI is now a material part of the risk and compliance technology stack. The 300 financial services firms surveyed* forecast \$703.7 million of AI spend in 2026.

Compared with the \$2.37 billion of total RegTech spend reported by a similarly profiled sample in our wider 2026 research, this provides a directional benchmark of approximately 30% of risk and compliance technology budgets.

Unlike the headlines of 12 months ago, these investments appear to be paying off. All 300 institutions report some use or active exploration of AI. Across the seven domains assessed, 65.6% of reported AI activity is at production stage or beyond, while 58.4% involves AI taking action, rather than only providing information or recommendations. More than a quarter of institutions (27.7%) report realised returns above 50%, compared with 13.3% in the 1%–10% band. These are self-reported estimates rather than audited profitability measures, but the direction of travel is clear.

This improvement has coincided with the emergence of hybrid AI architectures. Almost four in five vendor respondents** use at least three different model types, and 69% include agents in their solutions. Agents already inform or materially influence decisions at 51% of institutions. However, the newest models do not appear to be driving returns by themselves: high-ROI firms differ most from laggards in the responsibility they give to rules-based AI and predictive machine learning (ML).

For this report, ROI leaders are institutions reporting realised returns of 51% or more; ROI laggards report returns of 1%–10%. Five lessons emerge:



1. Deploy beyond isolated use cases. ROI leaders score 4.9 points higher on AI Deployment Intensity, which combines deployment depth, autonomy and decision impact.



2. Target business outcomes, not only efficiency. Leaders are much more likely to report faster onboarding and approvals (+15.9 percentage points), faster product launches or market expansion (+15.6 points), and lower technology costs (+12.6 points).



3. Connect knowledge to action. Leaders over-index most strongly on information retrieval (+20.2 points), structured text generation (+18.2 points) and workflow orchestration (+17.5 points): finding the right evidence, turning it into a usable output and moving it through a process.



4. Use the right AI for the job. ROI leaders differ most from laggards in the responsibility they give to rules-based AI and predictive ML. They appear to be combining newer AI with proven, bounded technologies rather than simply chasing the latest models.



5. Solve the implementation layer. Leaders are less constrained by integration and regulatory uncertainty. Their challenges have moved downstream towards staff adoption, vendor limitations and translating policy into operational controls.

Spending alone is not a differentiator. Firms in the top quintile for AI spend per employee are no more likely to be ROI leaders than the rest of the market: 26.7% versus 27.9%. Investment matters, but value depends on how effectively AI is integrated into decisions, workflows and the wider operating model.

The gains are also not yet fully secured. According to our measures of AI Deployment Intensity and AI Control Maturity, 69.3% of institutions have deployment running ahead of their measured control environment. The average institution has only 4.4 of the 13 controls assessed.

This is therefore not a permanent division between winners and losers. Some low-ROI firms may be prioritising benefits such as improved detection, resilience or avoided losses that are harder to quantify in near-term returns. Some early leaders may be accumulating governance and assurance debt. Durable advantage will belong to institutions that preserve deployment momentum while raising controls in proportion to the autonomy, impact and risk of each use case.

*See [Appendix 1](#) for financial services survey scope

**See [Appendix 2](#) for vendor survey scope

The **current state of AI** in risk and compliance

Commercials

Return on investment

Last year's report argued that, despite certain headlines, generative AI did not have a fundamental ROI problem. Rather, many firms were deploying low-risk, low-value use cases while still absorbing the technology's high costs. The market's subsequent focus on measurable outcomes appears to have changed the picture.

Every respondent* reports a positive realised return, although most returns remain modest. More than a quarter of firms now report returns above 50%, while only 13.3% sit in the 1%–10% band.

This should not be read as an audited profitability measure. Almost one-third of institutions still identify difficulty measuring ROI as a deployment barrier, and responses are treated as estimates rather than verified financial returns.

“For teams suffering from AI fatigue, the metrics that matter [for marketing compliance] are operational: first-time approval, approval time and SLA compliance—alongside false positives and false negatives for regulatory safety.”

Joe Jordan, Co-founder, Adclear


“ROI is downstream of visibility. You cannot reliably answer the ROI question until you can see how AI is influencing the workflow.”

Matt Holmes, Chair, Responsible AI Working Group,
UKAI Trade Association

“AI has taken my review of a due-diligence data room from four days to four hours — and it can help me see the data more holistically rather than through individual compliance silos.”

Chief Compliance Officer, US wealth management firm

Figure 1: Realised return on AI investment



*Unless otherwise stated, all figures and analyses are based on a survey of 300 senior risk and compliance professionals in financial services. “Firms” refers to the institutions represented by those respondents. See Appendix 1 for the sample profile.

Leaders and laggards

For the rest of this report, “ROI leaders” are firms reporting realised returns of 51% or more. “ROI laggards” are firms reporting returns of 1%–10%. These thresholds are analytically useful groups, but they do not imply that low reported ROI necessarily reflects weak *long-term* performance.

ROI leadership is most common in the UAE (40%) and least common in France (15%). Asset managers (35%) and insurers (33%) report the highest incidence of ROI leadership among the main subsectors. FinTechs and incumbents are effectively level. The 20,000+ employee group reports the strongest results (54% ROI leaders). Firms with 500–999 employees also perform strongly, at 36%.

Despite the variation, these results are far more positive than those from 12 months ago. However, the full story is more alarming than it might seem. The remainder of the report explores firms’ AI status in more detail and seeks to identify key differences between the ROI leaders and the rest. Spoiler: firms have indeed traded risk for reward, as our previous paper recommended, but the speed of AI progress has not always been matched by governance.

“Regardless of how firms choose to develop GenAI, the message for industry leaders is simple: treat risk not as a brake but as a design variable. Risk is the currency that we trade for reward, but it should be paid for with rigor to avoid reputational and regulatory costs.”

Generative AI in Risk & Compliance 2025

Realised benefits

Table 1: Commercial benefits by realised ROI

% of respondents

	ROI Group				
Commercial benefit	Laggards 1%–10%	11%–25%	26%–50%	Leaders (51%+)	Leaders vs Laggards Gap
Reduced manual effort	47.5%	50.5%	51.2%	50.6%	3.1%
Increased productivity or throughput	37.5%	49.5%	39.0%	42.2%	4.7%
Faster onboarding, approvals or reviews	27.5%	40.0%	42.7%	43.4%	15.9%
Improved detection of risk, fraud or control failures	52.5%	32.6%	40.2%	33.7%	-18.8%
Faster regulatory reporting and evidence production	27.5%	38.9%	40.2%	28.9%	1.4%
Improved customer or employee experience	30.0%	30.5%	35.4%	30.1%	0.1%
Reduced technology costs	17.5%	24.2%	28.0%	30.1%	12.6%
Reduced losses, incidents or penalties	20.0%	15.8%	29.3%	28.9%	8.9%
Reduced headcount or outsourcing costs	30.0%	13.7%	28.0%	14.5%	-15.5%
Improved resilience and incident response	20.0%	18.9%	12.2%	20.5%	0.5%
Faster product launches or market expansion	2.5%	14.7%	17.1%	18.1%	15.6%

“It is now time for the industry to recognise risk and compliance as the strategic engine that it can be.”

Generative AI in Risk & Compliance 2025

ROI leaders are not simply automating the same compliance work more cheaply. They are more likely to use AI to remove friction from business activity: faster onboarding, approvals and reviews; faster product launches or market expansion; and lower technology costs. The common thread is that AI is being connected to revenue, customer journeys and operating leverage, rather than contained within a narrow efficiency programme.

“We are not trying to replace the compliance officer or their decisions. We are reducing manual effort so compliance teams can spend more time on strategy and decision-making, and move from being seen as a cost centre to becoming a business enabler.”

Pooja Shah, Director of Products & Client Experience, Sherlocq

ROI laggards are more likely to report reduced headcount or outsourcing costs and improved detection of risk, fraud or control failures. The first is easy to count but strategically capped. The second may be among the most valuable long-term outcomes in the study, yet it is difficult to translate into a short-term percentage return because the counterfactual loss, fine or incident is unobserved.

Some of the difference between leaders and laggards is therefore a difference in the accounting lens. Benefits that reach the P&L quickly are more likely to push firms into the leader group. Benefits that improve assurance, detection or resilience may create significant economic value without appearing immediately in an ROI calculation.

“Once you reach a certain maturity, this stops being only a compliance story; it becomes a speed, cost and revenue story.”

Ari Applbaum, Vice President Marketing, LeapXpert



Perception vs reality: The gap between vendors' perspectives and institutional results

Vendors view AI value primarily through the workflow: reduced manual effort and improved detection are each cited as key commercial drivers by more than 70%, substantially above institutions' own results. The two sides align on productivity, but vendors are less likely to identify broader outcomes such as lower technology costs, improved resilience, better user experience and faster market expansion.

This may reveal an opportunity for vendors. As buyers become **more focused on outcomes**, providers can move beyond selling efficiency or control improvement and position themselves around the broader business value of risk and compliance: enabling faster decisions, smoother customer journeys, lower operating friction and ultimately growth.

Spending

The 300 institutions in this study forecast \$703.7 million of AI spend within risk and compliance during 2026. In [The Global State of RegTech 2026](#), a similarly profiled sample of 300 institutions reported \$2.37 billion of total RegTech spend. This provides a directional benchmark of approximately 30% of reported risk and compliance technology spend being allocated to AI.

A closer look at the data shows that this figure varies significantly with firm size. Spend rises sharply with organisational scale. The step-change begins above 1,000 employees and accelerates among firms with 5,000 or more. This points to a meaningful fixed-cost threshold: scaled AI requires more than model access. It requires data preparation, integration, assurance, specialist skills, security, procurement and continuous monitoring. Larger institutions can spread those costs across more workflows and a wider control environment.



For the remainder of the report, we define “high spenders” as the top 20% of firms by forecast 2026 AI spend per estimated employee.

The largest firms do not allocate the largest proportion of their broader RegTech budget to AI. The share peaks in the 10,000–19,999 employee band, before falling among firms with 20,000 or more employees. That is consistent with the largest institutions maintaining broader and more expensive legacy, cyber, resilience and financial-crime stacks alongside their AI programmes.

Table 2: Forecast 2026 AI spend by firm size

Employee size	Mean AI spend (\$m)	Mean RegTech spend (\$m)*	AI spend as % of RegTech spend
50 - 249	\$0.08	\$1.80	4.37%
250 - 499	\$0.25	\$3.90	6.38%
500 - 999	\$0.43	\$6.10	7.08%
1,000 - 4,999	\$1.55	\$7.80	19.83%
5,000 - 9,999	\$4.72	\$14.00	33.72%
10,000 - 19,999	\$6.94	\$15.80	43.89%
20,000 +	\$11.00	\$42.50	25.89%

*Figures taken from The Global State of RegTech 2026 research initiative.

Spend is not a proxy for execution quality



ROI leaders spend materially more than ROI laggards, **but high spending alone does not predict leadership.**



Leaders spend around **36%** more per employee than laggards, yet firms in the top spending quintile are no more likely than the rest of the market to report ROI above **51%**.



FinTechs, for example, spend approximately **29% more per estimated employee** than incumbent institutions (\$645 versus \$501), but report almost identical rates of ROI leadership (27% versus 28%).

Deployment

Types of AI

This report is no longer limited to generative AI because real-world deployment in risk and compliance is hybrid. Firms are combining deterministic rules, predictive machine learning, large language models (LLMs), multimodal systems and agents. These techniques are not substitutes. They solve different problems and are increasingly assembled in sequence: rules constrain; predictive models score; LLMs interpret and communicate; multimodal systems expand the evidence base; and agents action or coordinate work across systems.


“Using AI” is no longer a sufficiently precise description. The relevant questions are which model is used, how deeply it is embedded, how independently it can act and how much influence it has over a regulated decision.

The longest-established model types are still entrusted with the most consequential work. Rules-based systems and predictive ML each record a Decision Impact* Score of 68.3, based on a weighted average of the responses in figure 2 scaled from 0–100, while agents score 55.2. This is rational, not necessarily conservative: deterministic logic and supervised models are narrower, more familiar and easier to test against expected outputs. What firms appear to trust is not determinism alone, but boundedness, testability and control.

The more flexible and general-purpose the technology, the further firms currently keep it from the final decision. LLMs, multimodal systems and agents can interpret a broader range of inputs and coordinate more complex tasks, but their variability also creates a wider control surface. They are therefore more commonly used to retrieve information, interpret evidence, explain results or prepare actions than to determine high-stakes outcomes directly.

Figure 2: How AI model types influence decisions

% of respondents



*The Decision Impact Score is a weighted average of survey responses shown in Figure 2, where 100 would mean all respondents say the model type materially influences outcomes and 0 would mean no decision impact.

Advanced models are nevertheless moving beyond preparation. For agents, 14.7% of respondents say the technology materially influences outcomes and a further 36.3% say it informs decisions. In other words, agents already shape decisions in some form at 51% of institutions, even though only a minority permit them to drive the outcome directly.



The strongest ROI differentiators are not the newest models. Compared with ROI laggards, leaders record the largest Decision Impact Score uplifts for rules-based AI (+9.4 points) and predictive ML (+8.6), versus +3.3 for agents.

The possibility that firms use multiple types of AI in sequence blurs the meaning of these results. It is not clear whether ROI is driven by any single model type or by that model used in sequence with others. It may, for example, be that rules-based AI models correlate with ROI leadership because firms are orchestrating them intelligently with language models and agentic layers. Or it may simply be that those firms are using rules-based models very well in isolation. We cannot conclude with certainty with the current data.

The supply side reinforces the shift towards hybrid AI. Almost four in five vendor respondents use at least three model types, while more than half use four or more. LLMs and rules-based systems are both present in around 85% of solutions, suggesting that probabilistic models are being added to deterministic foundations rather than replacing them. Agents are already used by 69%, but typically as part of a wider architecture.

“Before moving from human approval of every action to greater AI autonomy, an institution needs ‘situated evidence’: the system observed in the particular workflow, at a regular cadence, with its variation characterised.”

Reva Schwartz, Co-Founder, Civitaas; former Research Scientist, NIST

Table 3: AI model types, strengths, risks and decision impact

AI type	Description	Strengths	Risks	Decision Impact Score /100	Incidence in vendor solutions*
 Rules / symbolic AI	Executes explicit logic and predefined rules.	Deterministic decisions, obligations, thresholds, policy logic	Narrow usage, lacks deeper “awareness” of context	68.3	84.5%
 Predictive ML	Learns patterns from historical data to score, classify or forecast.	Scoring, detection, forecasting, anomaly identification	Explainability, drift, bias	68.3	63.1%
 LLMs	Generates or retrieves natural-language content from broad or controlled knowledge.	Search, summarisation, drafting, knowledge retrieval	Hallucination, data leakage, weak structured-data reasoning, token cost	62.6	85.7%
 Multimodal AI	Interprets and combines text, image, audio and video.	Document, voice, image, video and evidence interpretation	Authenticity, robustness, provenance	57.3	50.0%
 AI Agents	Coordinates multi-step tasks and actions across tools with limited supervision.	Workflow orchestration, task execution, monitoring, remediation	Identity, permissions, accountability, runaway execution, token cost	55.2	69.0%

*taken from this report’s vendor survey (n=100) ; see [Appendix 2](#)

Depth and autonomy

Deployment depth measures how far AI has progressed from exploration and pilots into multiple production deployments and core operational processes. Autonomy measures the highest level of independent action, from providing context to executing within predefined automated controls. Both are scored from 0–100 across seven risk and compliance domains.



Domain scores are averages across respondents, where 100 would mean that all respondents reported the maximum level of deployment depth or autonomy.


Depth and autonomy move closely together at firm level, but the exceptions are revealing. Cybersecurity, financial crime compliance and compliance management sit at the top of the chart. Data and technology risk are relatively deeply deployed but less autonomous, consistent with the broad consequences of acting incorrectly across systems, models, third parties and enterprise data.

Conduct Risk remains lower on both measures because surveillance and conduct decisions often depend on intent, context and a defensible evidential chain. ESG remains the least developed domain, reflecting weaker data foundations, less standardised decision processes and less consistent regulatory pressure.



Vendors appear more conservative than institutions about how far AI has moved into production. In Financial Crime and Compliance Management, the average vendor estimate of institutions' progress with AI is around half a stage below institutions' self-reported position. The gap is smaller on autonomy: vendors place institutions only around 0.2–0.3 stages lower. Most believe firms are comfortable with AI taking action subject to human approval (57.1%), but only 17.9% see appetite for review-by-exception or autonomous execution within predefined controls.

Figure 3: AI deployment depth versus autonomy



The pattern broadly follows the maturity of the underlying RegTech markets: Financial Crime and Information & Technology Security also led the RegTech Adoption Index in our previous study. Compliance Management is the notable exception. Its overall technology maturity (as measured by the RegTech Adoption Index) is lower, but text-heavy work such as regulatory research, interpretation, assessment and reporting is unusually well suited to LLMs and retrieval systems, which explains its higher scores in this study.

FinTechs and incumbent institutions have almost identical overall deployment scores. The differences appear within domains rather than between business models: incumbents are further ahead in financial crime and cybersecurity, while FinTechs show greater autonomy in conduct and technology-risk workflows. “Digital native” is therefore a weak proxy for AI deployment across the entire control environment.

The AI Deployment Intensity Index



The 2026 Global AI Deployment Intensity Index: 51.9/100

Understanding the index

(depth + autonomy + decision impact) ÷ 3
= AI Deployment Intensity

The index combines three equally weighted dimensions, each normalised to 0–100. It measures how deeply and consequentially AI is used. It is not a quality score and does not assume that more is always better.

Deployment depth is the average stage reached across seven domains, from exploring use cases to embedding AI in core processes. Deployment autonomy is the average highest level of action across the same domains, from providing context to acting autonomously within predefined controls. Decision impact captures the role of five model types, from preparatory support to materially influencing outcomes.

The apparent geographic convergence masks a much larger firm-level divide. The top 20% of institutions average 70.2, while the bottom 20% average 33.7. This divide has commercial consequences: high-deployment firms report a higher average ROI midpoint (52.4% versus 35.6%). It seems that, for now, deeper, more autonomous and impactful AI deployments are paying off.

Figure 4: Global AI Deployment Intensity by jurisdiction




Table 4: AI Deployment Intensity components

Segment	Deployment Depth	Deployment Autonomy	Decision impact
Global	49.2	44.2	62.4
UK	48.8	44.7	62.8
France	50.4	45.5	63.7
US	49.9	43.9	62.5
Australia	51.5	45.6	65.1
Singapore	45.7	41.1	60.4
UAE	48.1	44.0	59.5

Capabilities

Table 5: AI capabilities currently deployed

Capability Group	Capability	Description	% firms deployed in 2026	% firms deployed or planning deployment in 2025
Information Handling	Information Summarisation & Comparison	Condensing or comparing documents/data	55.0%	55.4%
Information Handling	Information Retrieval	Finding relevant content from internal/external sources	52.7%	49.2%
Text Interaction & Composition	Structured Text Generation	Drafting/editing reports, templates, or responses	48.3%	40.0%
Text Interaction & Composition	Conversational Interaction	Chat-style or natural language interfaces	45.0%	41.5%
Engineering	Data Generation & Structuring	Synthetic data, labelling, organising or connecting structured data	40.7%	60.0%
Reasoning & judgement	Scenario Evaluation	Checking outcomes against rules or constraints	39.0%	30.8%
Reasoning & judgement	Hypothesis & Scenario Generation	Suggesting explanations, next steps, or analysis	34.3%	36.9%
Reasoning & judgement	Workflow Orchestration	Coordinating multi-step tasks, systems or agents	29.3%	N/A
Engineering	Coding & System Logic	Writing or updating code, scripts or system logic	24.3%	53.8%
Media Generation & Analysis	Media Interpretation & Summarisation	Analysing media like scanned documents, diagrams, digital graphics or videos	19.7%	44.6%
Media Generation & Analysis	Multimodal Output Generation	Creating images, video, audio, charts, or slides	16.7%	41.5%

“Scenario evaluation at this level is impressive, but it is also a worry if the data quality and human oversight are not there.”

Katsuko Ishizeki-Chaudhari, Expert Advisor (ex-UK regulator and banking C-suite)

“A task that used to take me two days can now take about an hour with an AI agent, with me reviewing the output for mistakes.”

Ricardo Morais, Consultant, SET Advisory



2025 figures are from the financial-services subset of Generative AI in Risk & Compliance 2025 and combine deployed and planned use of generative AI specifically; they are directional comparators, not like-for-like year-on-year measures.

Production deployment has settled around information handling. Summarisation, retrieval, structured text and conversational interfaces are the most commonly deployed capabilities. These capabilities fit existing work, use familiar data and can remain assistive while still producing visible time savings.

The 2025 data suggests that not all plans or proofs of concept were taken forward into 2026. Structured text generation and scenario evaluation now exceed the previous deployed-or-planned benchmark. By contrast, coding, multimodal outputs, media interpretation and data structuring remain well below last year's combined deployment and intention figures. This does not prove that projects were cancelled, but it shows where ambition has, and has not, converted into production.

“The value in applying AI to regulation does not sit in the model, it sits in the corpus. Laws, guidance, enforcement actions and supervisory material need to be brought together across jurisdictions, curated and kept current. That underlying data is what gives the output context and makes it more defensible.”

Bhavin Shah, CEO & Founder, Sherlocq



ROI leaders are differentiated by their ability to connect knowledge and action. They are 20.2 percentage points more likely than laggards to use information retrieval, 18.2 points more likely to use structured text generation and 17.5 points more likely to use workflow orchestration. Summarisation is also common, but the stronger commercial pattern appears when firms retrieve the right evidence, turn it into a usable output and route it through a workflow.

Access and procurement

The anticipated “SaaSocalypse” is not visible in these results. Institutions are combining general-purpose models, enterprise software, specialist RegTech, custom applications and in-house models. In-house and open-source development is the least common route, despite the falling cost of prototyping.

“A raw foundation-model answer is not enough for a regulated workflow; you need citations, an audit trail, permissions and human validation around it.”

Ian Howard, Director, Neotas

That is consistent with the build-versus-buy pattern in our wider RegTech research. AI lowers the cost of creating a proof of concept; it does not remove the burden of production engineering, data governance, security, validation, maintenance and regulatory evidence. Specialist vendors and enterprise providers can spread those obligations across customers and keep products aligned with fast-changing model capabilities.

“You can build some very clever interfaces quickly, but in this space the quality of the data is what really matters.”


Lauren Gomersall, Solutions Engineer, CUBE

The prevalence of granting general access to foundation models can reduce unsanctioned “shadow AI” by giving employees an approved route, but it places more responsibility on identity, permissions, data classification, training and user judgement.

“Shadow AI is real. We can pretend that putting a policy in place has stopped it, but the technology is too accessible and people are too interested in using it. Risk managers are better served by acknowledging that behaviour and bringing it inside controlled systems.”

Kathryn Carlson, Chief Product Officer, Riskconnect

Figure 5: AI access and procurement methods



“If employees are going to use foundation models, they need enterprise access with guardrails; you do not want client data being put into public models outside the controlled environment.”

Chief Compliance Officer, US wealth management firm

“Giving everyone a licence and letting them organically figure out how to use AI is an incredibly dangerous deployment model.”

Matt Holmes, Chair, Responsible AI Working Group, UKAI Trade Association

Barriers


The cost of building technology internally and staff resistance have fallen sharply in prominence. Data quality and availability (51.7%), lack of internal expertise (49.7%) and regulatory uncertainty (45.3%) now lead. These are the relevant constraints when pilots move to production: firms need dependable inputs, people who understand both AI and regulation, and defensible interpretations of what supervisors expect.

The ROI story has also matured. Uncertainty over whether AI can create value has receded, but difficulty measuring the exact return has become more prominent. Firms are now trying to attribute benefits across avoided losses, faster decisions, staff capacity, customer outcomes and technology costs, many of which do not sit neatly in one budget or reporting line.



ROI leaders are not operating in lower-friction environments. But their reported challenges are different. They are less constrained by integration and regulatory uncertainty, but more likely to encounter staff resistance, vendor limitations and the practical challenge of turning policy into operational controls.

Figure 6: Barriers to deployment



Risk

Intensity

AI creates risk in two directions. Risk and compliance teams must govern the technology used by their own institutions, while also responding to AI-enabled fraud, cyber threats, manipulation and other external harms. The greatest current pressure is concentrated in technical domains, where AI expands the attack surface, changes data flows and creates new model and third-party dependencies.

Vendors and institutions agree that AI risk is concentrated in Cybersecurity and Data Risk, and their scores across the main technical domains are remarkably close. This hierarchy reflects where AI changes the underlying exposure most directly: access, software, data provenance, model behaviour, third parties and operational resilience.

Vendors are less likely, however, to see material risk in Conduct Risk and Compliance Management. This may be down to different vantage points. Technology providers are close to model, data and system risk, while institutions carry the wider evidential, conduct and regulatory-accountability consequences once AI is embedded in business decisions.

“The model is only part of the risk; the setting and workflow it sits inside can matter even more to the decisions people make.”

Matt Holmes, Chair, Responsible AI Working Group,
UKAI Trade Association

Table 6: AI risk intensity

% of respondents

Domain	Institutional Perspective							Vendor Perspective
	0: No risk	1: Low	2: Moderate	3: High	4: Very high	5: Critical	Mean risk /5	Mean risk /5
Financial Crime Compliance	2.3%	18.0%	27.3%	32.0%	13.7%	6.7%	2.57	2.45
Conduct Risk	4.7%	13.3%	27.3%	29.3%	19.0%	6.3%	2.64	2.32
Compliance Management	3.0%	13.0%	27.0%	31.7%	17.3%	8.0%	2.71	2.32
Cybersecurity	0.0%	9.7%	23.3%	29.3%	26.3%	11.3%	3.06	3.19
Technology Risk	0.7%	14.0%	25.3%	34.7%	18.7%	6.7%	2.77	2.69
Data Risk	2.3%	10.3%	20.0%	26.0%	29.3%	12.0%	3.06	2.98
ESG	10.7%	27.3%	29.7%	18.0%	9.3%	5.0%	2.03	1.75



The most commercially successful and most deeply deployed firms perceive more risk, not less. ROI leaders rate an average of 4.5 of the seven domains as presenting high, very high or critical AI-related risk, compared with 3.3 among laggards. High-deployment firms rate 5.0 domains at these levels, compared with 2.8 among low-deployment firms, and report higher risk in every domain.

This is not evidence that deployment has failed. It suggests that exposure and visibility grow together. Firms that use AI more consequentially encounter more edge cases, dependencies and control questions. Lower reported risk can indicate strong management; it can also indicate more limited use or awareness.

Drivers

The most widely selected challenge is external: the use of AI to conduct fraud and scams, cited by 39.3% of firms. AI-enabled cyber weapons, prompt injection and adversarial manipulation also rank highly. These threats sit directly within the remit of financial crime, cyber and technology-risk teams, making AI both a governed technology and an accelerant of the risks those functions already manage.

Governance and operating issues account for 41.3% of all top-five selections, compared with 33.9% for model and technical risks. Inadequate human oversight, unclear accountability, shadow AI, missing inventories, regulatory non-compliance and weak traceability are therefore a larger combined problem than any single model failure mode.


Vendor respondents reinforce the internal nature of the challenge. Governance is most often ranked as the leading risk category, by 41.7%, while model risk is most commonly placed second.

The risk agenda has broadened beyond the early public debate around hallucination and bias. Each was selected by 21.7% of respondents, still material, but less frequently ranked among the top five challenges than human oversight, data leakage and AI-enabled fraud. Institutions are still concerned about whether models produce reliable and unbiased outputs, but the findings point to a wider question: whether unreliable outputs can be detected, what actions AI is permitted to trigger, and who is accountable for the resulting decisions. This extends the focus beyond model outputs to the wider control environment. It requires institutions to know where AI is used, who owns it, what it can access, how its outputs are reviewed, and whether its use and resulting actions can be reconstructed and evidenced.

“We are at a tipping point. Autonomous agents put the risks of both over-suppression and non-compliance on steroids. The focus now has to be on putting guardrails around traditional, live and autonomous engagement.”

Neal Keene, CTO, Gryphon

Figure 7: Top AI risks



% of firms ranking each risk or threat among their top five AI challenges.

■ External / adversarial ■ Model / technical ■ Governance / operating

Governance and controls

The four most common controls tell a coherent story: firms have written policies, retained human oversight, trained staff and begun to capture evidence. These are necessary foundations for establishing intent, accountability, and a minimum audit trail.

“I am genuinely surprised by how few firms appear to have these controls in place; you would hope the numbers were much closer to at least 75% or 80%.”

Chief Compliance Officer, US wealth management firm

“If governance means ‘put a policy in place and that will do’, then governance is not being understood.”

Katsuko Ishizeki-Chaudhari, Expert Advisor
(ex-UK regulator and banking C-suite)

“AI can generate an answer, but compliance teams need a defensible decision. The challenge is building that defensible layer around the AI’s output.”


Mark Holmes, Head of Solutions Consulting, AscentAI

However, the typical control environment is policy-forward but operationally thin. Only one-third of firms maintain an AI inventory or use-case register. Even fewer have processes for risk classification (31.3%), pre-deployment review and approvals (29%) or monitoring and performance management (26.7%).

These omissions leave critical gaps. Without an inventory, firms cannot establish scope. Without monitoring, they cannot identify drift or changed behaviour. Without incident management, they cannot respond consistently. Without supplier diligence and assurance, they cannot evidence that outsourced systems meet the institution’s own obligations.

High-control firms are 30 percentage points more likely than low-control firms to use specialist AI-enabled risk and compliance solutions. The direction or existence of causality is unclear. But it may be that high Control Maturity leads firms somewhat inevitably toward dedicated, specialist solutions.

Figure 8: AI governance and control capabilities in place



They are also more likely to cite data, expertise and regulatory uncertainty as barriers to AI adoption than low-control firms. Low-control firms, by contrast, are far more likely to report insufficient governance frameworks and an inability to implement policies effectively. This is a progression curve: once the basic framework exists, the limiting problem becomes operationalising it across data, people and use cases.



High-deployment firms have more policies, oversight, testing and monitoring than low-deployment firms, but the uplift is not sufficient to match their greater intensity.

Control Maturity also does not translate immediately into a higher probability of being an ROI leader: 23.3% of both the highest- and lowest-control quintiles report returns above 50%. High-control firms do, however, use more capabilities and report a broader set of commercial benefits. Controls may be enabling breadth and durability before they show up in measurable ROI.

Firms relying exclusively on general-purpose access routes (General foundation-model access and/or Enterprise-software AI only) have materially thinner control environments. This result is driven primarily by firms providing general access to foundation models without also deploying specialist, custom or enterprise-embedded systems. The foundation-model-only subgroup averages fewer than three of the 13 measured controls, despite recording slightly above-average Deployment Intensity.

Lighter governance may be proportionate for specific low-risk, assistive use cases. In some cases, it can also be explained by the firm being a deployer, without visibility or access to underlying technical architecture or training data. Even so, use cases built on general purpose models still require appropriate policy, training, permissions, data controls, evidence capture and incident response. The data suggests that firms focused on general-purpose models are substantially behind even in these areas.

“If AI contributes to a client deliverable and something goes wrong, you need to be able to trace how it was used, what it produced and where the mistake came from.”

Ricardo Morais, Consultant, SET Advisory

“My main concern is around the quality of the evidence the organisation is relying on. Evaluation results about model behaviour on canned data do not represent how these systems operate in real settings and how that changes over time.”

Reva Schwartz, Co-Founder, Civitaas; former Research Scientist, NIST

The AI Control Maturity Index

The 2026 Global AI Control Maturity Index: 41.9/100

Understanding the index

(0.8 x control coverage) + (0.2 x implementation confidence) = AI Control Maturity

Control coverage is the share of 13 governance and control capabilities in place. Implementation confidence is a narrow respondent-level adjustment based on whether firms report an insufficient framework and/or an inability to implement controls effectively.

The weighting is deliberately simple. Observable control coverage carries 80% of the score. Implementation confidence carries 20%, preventing a firm that reports widespread controls but also fundamental implementation failure from being treated as fully equivalent. The confidence component is perception-based and should not be mistaken for independent validation.


The index is a comparative measure. It asks how developed the firm's control environment is relative to the survey framework. It does not determine whether a specific use case complies with applicable law or whether every implemented control is effective.

The global score of 41.9 is driven by average coverage of only 33.9%, equivalent to approximately 4.4 of the 13 controls. Jurisdictional averages are again compressed, and firm-level scores range much more widely, from approximately 6 to 82.

Deployment Intensity and Control Maturity are positively related, but only weakly. Firms that deploy more tend to add controls, yet the relationship is far from automatic. This is the basis for the report's central metric: the distance between the intensity of use and the maturity of the supporting control environment.

Figure 9: Global AI Control Maturity by jurisdiction

Maturity Index Scores



While ROI leaders have higher average Control Maturity, the difference is slight, at 1.8 points, compared with the 4.9-point difference in Deployment Intensity. This suggests that returns are currently being created by more consequential deployment faster than they are being secured by stronger controls. A simple multivariate check supports that interpretation: after accounting for country, firm size and subsector, Deployment Intensity retains a positive association with self-reported ROI, while Control Maturity does not. **The sample size limits the significance of these results, but they are at least directionally consistent.**

Table 7: AI Control Maturity components

Segment	Control coverage	Implementation confidence
Global	33.9	73.8
UK	31.2	73.3
France	32.9	67.5
US	34.3	78.5
Australia	34.6	75.0
Singapore	31.3	73.3
UAE	39.2	68.8

The AI Readiness Gap



The 2026 Global AI Readiness Gap: +10

Understanding the gap

Deployment Intensity - Control Maturity = AI Readiness Gap


A positive score means deployment is ahead of the measured control environment. A negative score means controls are ahead of deployment. The gap is calculated from unrounded respondent-level scores, so displayed rounded components may not subtract exactly.

The gap is not a finding of non-compliance and does not imply that every positive-gap use case is unsafe. It is a directional measure of imbalance. A larger positive score means the institution is using AI more deeply, autonomously and consequentially than its governance and control coverage would suggest.

The global gap is +10, and 69.3% of firms have a positive score (have deployments that appear to have outgrown their control environment). France records the largest country-level gap (+13.4), followed by the UK (+12.5) and Australia (+11.4). The UAE records the smallest (+5.4), supported by the strongest measured control coverage in the sample.

At firm level, the top gap quintile averages 63.8 on Deployment Intensity and only 31.5 on Control Maturity, producing a gap of +32.3. These firms perceive more risk and are more likely to report ROI above 50% than firms with zero or negative gaps. Financial returns can therefore arrive before broader organisational readiness.

Figure 10: AI Readiness Gaps by jurisdiction



That does not make the gap economically harmless. Firms with positive gaps report fewer types of realised commercial benefit than firms whose controls are level with or ahead of deployment. The emerging pattern is one of narrower, faster value capture accompanied by assurance debt: today's measurable returns may be greater, but their repeatability and resilience remain less certain.

Closing the gap may not mean slowing every deployment until every control is perfect. It means raising the control curve without losing useful momentum: prioritising inventories, risk classification, pre-deployment review, monitoring, incident response and evidence according to the autonomy and decision impact of each use case.

“The limiting factor is not engineering, use cases or budget; it is the ability to prove to the regulator that as agentic deployment accelerates, the right oversight is in place to assure safe and appropriate consumer interaction and outcome.”

Kent Mackenzie, Chief Commercial Officer, Aveni

For now, the industry may have solved its ROI problem. But in doing so, there is evidence that it could have created a governance gap.

“The bank that solves the oversight problem will be the bank that wins the race to become the most agentic bank.”

Kent Mackenzie, Chief Commercial Officer, Aveni



CYBORG ECONOMICS

Expert Economics, Supercharged by AI

Cyborg Economics: the value at stake in AI readiness

For this research, [Cyborg Economics](#) estimates that **incomplete conversion of GenAI's potential into realised performance leaves around \$1 trillion of annual productivity-equivalent value at stake. Its central modelled estimate is \$0.90 trillion, with a wider sensitivity range of approximately \$0.60 trillion–\$1.21 trillion.**

Cyborg Economics constructed the estimate using a bottom-up approach: combining global economic and labour-income data, the [ILO's measure](#) of occupational exposure to GenAI, and published evidence on productivity gains from AI-assisted work. It then estimates how much of that potential may remain unrealised when adoption does not translate fully into scaled, measurable performance.

In practice, incomplete conversion can reflect weak data foundations, unclear ownership, immature governance, poor workflow integration or difficulty scaling successful applications. The estimate measures gross labour-

productivity-equivalent value at stake, rather than net corporate profit or the economic value of closing any single governance metric. The estimate also focuses on GenAI, which is not the totality of enterprise AI.

The estimate is therefore an order-of-magnitude scenario, not a forecast or a direct monetary value for the Parker & Lawrence AI Readiness Gap.

Nicholas Herrick, founder of Cyborg Economics, is a former senior UK government economist.

He led the Department for Transport's Economics Centre of Excellence, advised the Cabinet Office and HM Treasury, and delivered appraisals on major UK infrastructure projects. His current work focuses on the economics of AI adoption, enterprise readiness, productivity and value realisation.

The role of regulation

The AI Readiness Gap is, to some extent, structural. AI capabilities are advancing faster than governance frameworks can mature, while institutions face strong **commercial incentives** to deploy them. Regulators, meanwhile, are still defining how best to oversee a technology evolving faster than traditional rulemaking cycles.

AI regulation differs substantially in legal force, but much less in its underlying expectations. Across the six jurisdictions in focus, regulators generally avoid prescribing a particular model architecture or development process. Instead, they focus on the risk of the use case, the decisions affected and the quality of the governance and evidence surrounding the system.

The EU AI Act is the most binding and detailed regime in scope. Yet even its requirements for high-risk systems centre on risk management, data quality, logging, documentation, human oversight, robustness, cybersecurity and accuracy, without reaching into technical design.

Principles-based does not mean permissive. It shifts more of the interpretive burden to firms. Risk and compliance teams must decide what “reasonable”, “proportionate”, “fair”, “secure” and “explainable” mean for each use case, and retain evidence that those judgements remain valid as the system changes.

“A principles-based framework turns into a checklist when institutions treat principles as the finish line instead of the starting point. The best starting point is to get the organization to understand that risk is everyone’s job.”

Reva Schwartz, Co-Founder, Civitaas; former Research Scientist, NIST

51%

of financial institutions want clearer guidance from regulators on acceptable technologies in risk and compliance. [The Global State of RegTech 2026](#)

The regulatory core

Terminology varies, but eight control themes recur across the EU, UK, Singapore, US, Australia and UAE frameworks. The object of scrutiny is increasingly the use case and its evidence trail, rather than the model in isolation.

“Anything that may be audited needs to be repeatable. You do not want different answers each time, and in many cases the output also needs to be stored on the underlying record.”

Kim Alderman, Director, AI Innovation Lab, Riskconnect

Table 8: Common principles across the six regulatory environments

Common principle	The question firms must answer	Evidence in practice
Proportionality and use-case risk	What is the system’s intended purpose, and what harm could follow?	AI inventory, materiality tier, impact assessment and approval threshold
Accountability and governance	Who owns the model, decision and control?	Named accountable owner, defined roles, board reporting and escalation routes
Data, privacy and provenance	Are inputs lawful, representative and traceable?	Data lineage, legal basis, quality testing, access controls and source documentation
Fairness and outcomes	Do outcomes differ unfairly across customers or groups?	Bias testing, segment-level outcomes, threshold review and consumer-impact monitoring
Transparency and explainability	Can users, reviewers and supervisors understand the system’s role and output?	Model or system documentation, reason codes, disclosures and decision lineage
Robustness, security and resilience	Does the system remain safe under change, attack or failure?	Independent validation, stress testing, red teaming, fallback arrangements and incident plans
Human oversight and redress	Can a person intervene, override or challenge a decision?	Review triggers, stop controls, override rights and complaints or appeal routes
Lifecycle and third-party control	Is risk managed after deployment and across suppliers?	Logging, drift monitoring, revalidation, vendor due diligence and audit rights

The impact of regulation

The six-country results do not support a simple trade-off between regulation and AI deployment. To examine the relationship, we separate Regulatory Prescription (the extent to which AI-specific obligations are binding and detailed) from Regulatory Maturity (the clarity, completeness and operationalisation of the wider regime)*. They are distinct: a principles-based regime can be mature, while a prescriptive regime may still be difficult to interpret or implement.

The table places both measures alongside firms' AI Deployment Intensity, AI Control Maturity and resulting Readiness Gap. With only six jurisdictions, the comparison is diagnostic rather than causal.

France combines the clearest and most prescriptive regime with the second-highest Deployment Intensity. Only 30% of French firms cite regulatory uncertainty as a barrier, the lowest share in the sample. Yet France also has the largest Readiness Gap. Its control coverage is broadly comparable with peers; the difference is driven more by lower implementation confidence. A higher regulatory bar may make firms more critical of what “ready” means.

The UAE has the lowest Regulatory Maturity score in our assessment, and 52.5% of firms cite regulatory uncertainty. Yet it has the highest control coverage and the smallest Readiness Gap. Firms appear to be implementing controls without waiting for full system-wide harmonisation.

More broadly, there is no clear country-level relationship between Regulatory Maturity and deployment. The better question is therefore not whether regulation is “good” or “bad” for deployment. It is whether regulation gives firms enough clarity to take appropriate risk, while preserving flexibility as technology changes. Mature regulation can reduce uncertainty and support scale, but it can also raise the standard against which firms judge themselves. Less mature regimes may permit faster experimentation, while leaving more interpretation and harmonisation risk with institutions.

Across all six jurisdictions, the responsibility ultimately returns to the firm. Regulators may avoid prescribing the model, but they do not remove accountability for the outcome. The practical unit of compliance is the use case: what it can do, what can go wrong, who is accountable and what evidence proves that the controls are proportionate and effective.

*These are analyst-generated metrics; see [Appendix 3](#) for detailed definitions and methodology.

Table 9: Regulation and AI readiness by jurisdiction

Rank	Jurisdiction	Current regulatory model	Prescription /5	Regulatory Maturity /100	AI Deployment Intensity	AI Control Maturity	AI Readiness Gap
1	France	EU AI Act and GDPR overlay, with ACPR responsible for high-risk financial-sector AI. Key high-risk financial-use obligations apply from December 2027.	5.0	91	53.2	39.8	13.4
2	Singapore	MAS and IMDA frameworks supported by the MindForge operational toolkit and SAFR runtime safeguards for AI agents. Highly implementation-led.	2.0	85	49.1	39.7	9.4
3	UK	Existing regulatory frameworks applied through an outcomes-based approach. The FCA does not currently plan additional AI rules and is using live testing to develop practical expectations.	2.0	82	52.1	39.6	12.5
4	US	Voluntary NIST framework, revised interagency model-risk guidance and binding state-level statutes. Technically sophisticated but institutionally fragmented.	3.0	72	52.1	43.1	8.9
5	Australia	Existing general and sectoral laws, voluntary national guardrails and increasingly active ASIC and APRA supervision.	2.0	70	54.1	42.7	11.4
6	UAE	Detailed but flexible CBUAE guidance, alongside separate DIFC and ADGM supervisory approaches. Financial-sector expectations are strengthening, but system-wide harmonisation is still developing.	2.5	69	50.5	45.1	5.4

Building the AI stack in risk and compliance

Building the AI stack in risk and compliance

Financial Crime Compliance

Financial Crime Compliance

Category Deep Dive

	Model Applicability Score					
Subcategory*	Rules-based	Predictive ML	LLMs	Multi-modal	Agents	What this means in practice
Identity Verification	3	3	1	3	2	ML and multimodal AI are central to modern identity verification, supporting document authentication, facial matching, liveness detection and identification of suspicious patterns. Rules remain important for deterministic verification requirements, while agents can increasingly orchestrate checks and exception handling across multiple data sources.
KYC / KYB	3	2	3	2	3	LLMs and agents can significantly reduce the manual research involved in onboarding and periodic reviews. LLMs can extract information from documents, analyse adverse media and summarise complex ownership structures, while agents can gather data, perform checks and assemble cases for human review. Rules remain essential for mandatory checks and escalation criteria.
Fraud Detection & Prevention	3	3	2	3	2	Predictive ML remains one of the strongest technologies for fraud detection, identifying behavioural anomalies and patterns across large volumes of real-time data. Multimodal AI is increasingly important for document, biometric, voice and deepfake fraud, while LLMs and agents can enrich alerts, support investigations and coordinate response actions.
Transaction Monitoring	3	3	2	1	2	Rules and predictive ML remain the core detection technologies, with ML improving behavioural monitoring, anomaly detection and alert prioritisation. LLMs add value downstream by summarising activity and supporting investigations, while agents have strong potential to automate alert triage, gather supporting evidence and progress cases through investigation workflows.
Sanctions & Screening	3	3	2	1	2	Rules remain critical for screening against prescribed lists, while ML can improve matching, entity resolution and false-positive reduction. LLMs can help investigators interpret contextual information and complex ownership or relationship data, with agents increasingly able to conduct follow-up research and orchestrate screening investigations.
Financial Crime Risk Assessment	3	2	3	1	3	GenAI is particularly useful for bringing together fragmented risk information and turning it into explainable assessments. LLMs can synthesise risk factors, typologies, regulatory requirements and control evidence, while agents can gather inputs and continuously refresh assessments. Rules provide the underlying scoring methodology, with ML supporting more dynamic risk segmentation and prioritisation.

3 — **Strong fit:** One of the most useful technologies for this use case; proven and deployable today.

2 — **Useful:** Adds meaningful value, usually alongside other approaches rather than as the core technology.

1 — **Emerging / niche:** Relevant in some scenarios, but limited adoption or a relatively narrow use case today.

0 — **Limited fit:** Little practical value for this use case.

*Check [Appendix 4](#) for full taxonomy

** Parker & Lawrence Research analyst assessment. Scores represent practical model fit, not surveyed adoption, product performance or a vendor ranking.

AI Adoption and Risk

AI Deployment: % of firms reporting each level of deployment in Financial Crime Compliance

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
AI is embedded in core operational processes	12%	10%	15%	10%	20%	3%	15%
Multiple production deployments across the domain	29%	33%	28%	29%	23%	43%	18%
Limited production deployment in specific workflows	32%	23%	30%	35%	43%	20%	40%
Pilot or proof-of-concept deployments	19%	32%	23%	13%	13%	20%	15%
Exploring/evaluating use cases	8%	2%	5%	13%	-	13%	13%
No current AI activity	-	-	-	-	-	-	-

41%

of firms have deployed AI at scale in Financial Crime Compliance

52%

of firms report that AI is causing high levels of risk in Financial Crime Compliance

Financial Crime Compliance is one of the most mature AI categories, reflecting the long-standing use of rules-based and predictive models in screening and transaction monitoring. 73% report production deployment and 41% have reached multiple deployments or core-process integration. The next wave is increasingly agentic, with firms applying AI to automate multi-step workflows across KYB, screening and investigations, from gathering and validating information to resolving alerts and preparing cases for review.

AI Risk: % of firms reporting each level of risk in Financial Crime Compliance

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
5 - Critical AI-related risk 5	7%	2%	5%	14%	-	-	8%
4 - Very high AI-related risk 4	14%	17%	10%	12%	7%	20%	18%
3 - High AI-related risk 3	32%	38%	18%	33%	40%	30%	30%
2 - Moderate AI-related risk 2	27%	30%	40%	21%	27%	30%	25%
1 - Low AI-related risk 1	18%	12%	28%	18%	17%	17%	20%
0 - No identifiable AI-related risk 0	2%	2%	-	2%	10%	3%	-

Risk perceptions are relatively measured, with 52.3%* rating AI-related risk as high or above. Financial Crime teams have governed automated detection and decision support longer than most functions, and many AI-enabled threats are more scalable versions of familiar problems rather than entirely new ones. Deepfake documents and synthetic identities are increasing pressure on customer verification. Still, the more immediate challenge may be closer to home: maintaining evidence, explainability and investigator judgement as agents take on more of the screening and investigative workflow.

*Note: Percentages are rounded to whole numbers. Aggregated figures are calculated from unrounded data and may therefore differ from the sum of displayed percentages.

Building the AI stack in risk and compliance

Data Risk

Data Risk

Category Deep Dive

	Model Applicability Score					
Subcategory*	Rules-based	Predictive ML	LLMs	Multi-modal	Agents	What this means in practice
Data Governance	3	2	3	1	2	Rules remain central to ownership, classification standards, policies and governance workflows. ML is already used for automated classification and relationship discovery, while LLMs are increasingly helping users describe, search, curate and interpret data assets. Agents are emerging for stewardship and governance tasks, but are not yet the core operating model.
Data Lineage and Lifecycle Management	3	2	2	1	2	Lineage is still primarily driven by metadata, connectors and deterministic mapping of technical relationships. ML and LLMs can help infer undocumented relationships, enrich lineage and make complex flows easier to query or interpret. Agentic use is emerging around automated documentation and impact analysis, but remains secondary to the underlying lineage engine.
Data Quality	3	3	2	1	2	Rules and ML remain central to identifying invalid, incomplete or anomalous data, while LLMs can help generate quality rules, diagnose issues and explain likely root causes. Agents are emerging around remediation workflows, where they can investigate issues, coordinate fixes and re-test data quality after changes.
Data Protection	3	2	2	1	2	Rules remain fundamental for encryption, DLP, retention and protection policies, while ML is well suited to discovering sensitive data and identifying unusual access or exposure. LLMs can support investigation and policy interpretation, while agents are more likely to assist remediation than replace core protection controls.
Data Privacy	3	2	3	2	3	Privacy combines deterministic regulatory workflows with language-heavy analysis, making LLMs a strong fit for interpreting requirements, assessing processing activities and supporting reviews. ML can help discover and classify personal data, while agents can coordinate multi-step processes such as rights requests, assessments and evidence gathering.
Data Access and Usage Controls	3	2	2	1	1	Policy engines remain the core technology for access controls, masking and usage restrictions. ML can help classify sensitive data and identify unusual usage, while LLMs can simplify policy authoring and interpretation. Agentic use is still more relevant to provisioning and exception handling than enforcement itself.

3 — Strong fit: One of the most useful technologies for this use case; proven and deployable today.

2 — Useful: Adds meaningful value, usually alongside other approaches rather than as the core technology.

1 — Emerging / niche: Relevant in some scenarios, but limited adoption or a relatively narrow use case today.

0 — Limited fit: Little practical value for this use case.

*Check [Appendix 4](#) for full taxonomy

** Parker & Lawrence Research analyst assessment. Scores represent practical model fit, not surveyed adoption, product performance or a vendor ranking.

AI Adoption and Risk

AI Deployment: % of firms reporting each level of deployment in Data Risk

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
AI is embedded in core operational processes	8%	8%	8%	11%	10%	3%	5%
Multiple production deployments across the domain	26%	28%	18%	25%	27%	23%	33%
Limited production deployment in specific workflows	36%	30%	45%	36%	37%	43%	28%
Pilot or proof-of-concept deployments	21%	22%	20%	20%	23%	27%	15%
Exploring/evaluating use cases	10%	12%	10%	8%	3%	3%	20%
No current AI activity	-	-	-	-	-	-	-

AI Risk: % of firms reporting each level of risk in Data Risk

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
5 - Critical AI-related risk 5	12%	12%	13%	12%	13%	7%	15%
4 - Very high AI-related risk 4	29%	30%	35%	35%	23%	20%	20%
3 - High AI-related risk 3	26%	28%	28%	22%	20%	40%	25%
2 - Moderate AI-related risk 2	20%	18%	13%	24%	23%	17%	20%
1 - Low AI-related risk 1	10%	10%	13%	7%	10%	13%	15%
0 - No identifiable AI-related risk 0	2%	2%	-	-	10%	3%	5%

34%
of firms have deployed AI at scale in Data Risk

67%
of firms report that AI is causing high levels of risk in Data Risk

Data Risk has benefited from sustained investment in data governance, privacy and protection, helping 70% of firms move AI into production. Much of this adoption is arriving through established platforms, as mature vendors add AI capabilities to products already embedded within the enterprise. The market remains fragmented, however, and responsibility is often divided between data, privacy, security and technology teams operating different systems and control frameworks. Legacy estates, inconsistent metadata and disconnected ownership make integration difficult. This helps explain why only 8% have embedded AI into core processes: adoption is broad across the category, but still shallow across the operating model.

Data Risk produces the strongest risk signal in the survey, with 67%* rating it as high or above and 41%* as very high or critical. AI creates a fundamental tension for data teams: its value depends on making more organisational data accessible and reusable, while established governance is built around minimisation, purpose limitation and tightly controlled access. Enterprise adoption is also moving faster than many firms can resolve fragmented ownership, incomplete lineage and accumulated data debt. As AI becomes embedded through cloud platforms, enterprise software and third-party models, data can move beyond the systems and controls in which it was originally governed. The concern is therefore not only that AI may use poor or sensitive data, but that it expands the scale and speed at which existing weaknesses can propagate.

*Note: Percentages are rounded to whole numbers. Aggregated figures are calculated from unrounded data and may therefore differ from the sum of displayed percentages.

LeapXpert

Technology Leader: Data Governance



Company profile

LeapXpert provides a governed communications platform that brings employee-client conversations across WhatsApp, iMessage, WeChat, SMS, RCS, Telegram, Signal, and LINE into the enterprise data layer. The platform enables, governs and captures those interactions, then applies Communication Intelligence — agents, workflows and real-time data — to surface and act on compliance, cyber, relationship and commercial signals. Their agentic harness carries a firm's own AI agents into the live client conversation, under enterprise security, auditability and control.

Headquarters	New York, United States
Founded	2017
Founders	Dima Gutzeit , Avi Pardo
Website	leapxpert.com
Industry focus	Financial services (core), public sector and regulated or high-trust enterprises across legal, life sciences, energy, manufacturing, retail and hospitality.
Regional focus	Global; enterprise users in more than 45 countries.

Product differentiation

Governed communication data foundation: LeapXpert addresses the data-governance gap created when business conversations occur on consumer messaging channels or personal devices. Content, identity and metadata are captured in a controlled enterprise record that can feed archiving, surveillance, CRM, analytics and AI, while employees and clients continue using familiar channels.

Intelligence that acts: LeapXpert Intelligence both understands conversations and acts on them. Agents interpret what is happening and respond, at a level of autonomy the firm sets. Workflows carry out the steps that have to be right every time, such as identity checks, transactions, signatures and escalation. Real-time data gives firms governed access to what is happening as it happens, and the agentic harness carries their own agents into the client conversation, with identity, autonomy settings, human approval and full capture already in place.

Value drivers

-  **Increased productivity or throughput:** Live topic and signal views focus compliance, security and business teams on the conversations that warrant attention.
-  **Reduced manual effort:** Automated detection, alerts and integrations reduce manual archive searching, re-keying and hand-offs into downstream systems.
-  **Improved detection of risk, fraud or control failures:** Custom signals extend oversight across the governed message population, surfacing complaints, phishing, data leakage, MNPI and off-platform risks in context.
-  **Improved customer or employee experience:** Employees can use clients' preferred messaging channels through a governed professional identity, with relationship context available in existing workflows.

Key product features

Multi-channel enablement, capture & governance
Enables governed employee-client communication across WhatsApp, iMessage, SMS/RCS, WeChat, Signal, Telegram, LINE and collaboration interfaces including Microsoft Teams, Webex, and Slack. Messages, files, metadata, consent and identity are captured as enterprise data and can be exported to archiving and surveillance systems.

Live topic and trend intelligence
LeapXpert Signals analyses governed conversations in near real time, grouping them into topics and showing whether themes are emerging, stable or cooling. Users can filter by channel, organisational unit, role and time period to monitor compliance, cyber, client and business trends. Teams can also define what they want to detect in natural language, using a description, risk category, scope, examples, exclusions and optional keywords. Templates cover complaints, phishing, personal-data leakage, material non-public information and off-platform activity.

Agents in the client conversation
Firms bring their own AI agents; the agentic harness provides everything between that agent and a live client conversation. The agent works alongside a named relationship manager, at the level of autonomy the firm chooses: answering agreed questions, proposing a reply for approval, or passing the request straight to the person.

Workflows that complete the task
Workflows run the steps that have to happen the same way every time: verify who is on the other end, gather what is needed, complete the action and write the result to the record. A person can step in at any point, and every step is logged alongside the conversation itself.

AI applications

Model type	Capabilities	Key AI risks / threats mitigated
 Rules-based / symbolic AI Workflows define the steps that run the same way every time, for every client and every user. Governance and entitlement controls set which users, channels and data are in scope, how communications are retained and where a signal or completed step is routed. Nothing is generated on the fly, so behaviour is predictable and reviewable.	• Workflow orchestration: Workflows run the steps that have to be identical every time, from verifying identity to completing the action and writing the outcome to the record. • Scenario evaluation: Firm-set scope, channel, category and thresholds determine where a workflow or signal applies, whether a step proceeds or is held, and what follows.	• Unclear accountability or ownership for AI systems: The client defines the policies, the scope and what the platform is permitted to do. Workflows are released under the firm's control and can be paused at any time. • Inadequate human oversight of AI-driven decisions: A person can step in, take the conversation back or release a held step at any point, and every step is written to the audit log.
 LLMs Language models interpret plain-language signal definitions and message context, support multilingual matching, and generate the explanation shown with a detection.	• Information retrieval: Finds governed messages that match user-defined concepts and examples even where wording or language differs, then links the result to the source conversation. • Information summarisation & comparison: Summarises conversations, topics and trends and compares observed content with the signal definition, positive examples and exclusions. • Structured text generation: Produces the detection rationale and, in Maxen, conversation summaries, follow-up actions and suggested client-facing content.	• AI hallucinations/factual inaccuracy: Insights are grounded in captured enterprise conversations and can be checked against the source message. • Lack of transparency or interpretability: Users can see why a signal fired and drill into the underlying content rather than receiving an unexplained score. • AI leaking sensitive data: Tenant entitlements, PII redaction and governed capture restrict access.
 AI Agents Agents work in two places. Signals is itself agent-driven: agents monitor conversations against pre-built or client-defined trends, find the messages that matter, explain what they found and trigger alerts or workflows without being prompted. In the conversation itself, the agentic harness carries a firm's own agents alongside a named relationship manager, at the level of autonomy the firm sets, with every action captured and attributed.	• Information retrieval: Combines message content, metadata, scope, templates, examples and exclusions to identify relevant conversations across the governed data set. • Scenario evaluation: Assesses inbound communications against each trend definition and returns a match, non-match and reason; multilingual context allows one definition to operate across languages. Fast models also check each message against the firm's data-loss policies as it moves, and score deviations from established communication patterns in the Messaging Impersonation Detection module. • Workflow orchestration: Runs as messages arrive, raises email, API or webhook events, and can hand a request to a workflow or take one back.	• Inadequate human oversight of AI-driven decisions: The firm chooses the level of autonomy for each capability, from assisting the user to answering agreed requests directly, and the business logic and guardrails stay with the client. • Lack of traceability from AI outputs to business decisions: Each signal links to the source message, the reason and the audit record, and every agent action and workflow step appears in sequence in the same conversation history. • Prompt injection or adversarial manipulation: Consequential actions run as configured steps rather than generated ones, so an instruction hidden in a message cannot change what the process does.

Building the AI stack in risk and compliance

Cybersecurity

Cybersecurity

Category Deep Dive

	Model Applicability Score					
Subcategory*	Rules-based	Predictive ML	LLMs	Multi-modal	Agents	What this means in practice
Identity and Access Management	3	3	2	2	1	Rules still do most of the heavy lifting around entitlements, provisioning and policy enforcement. ML is well established for spotting unusual login or access behaviour, while LLM use is mostly around admin support and policy interpretation. Agentic infrastructure is still relatively peripheral to the core IAM use case.
Network Security	3	3	2	1	2	ML is already deeply embedded in network detection, particularly for anomaly detection and behavioural analysis. LLMs are increasingly being used to support investigation and explain complex telemetry, while agentic approaches are emerging around triage and response rather than replacing the core detection layer.
Endpoint Security	3	3	2	1	2	ML is a core part of modern endpoint detection, helping identify suspicious behaviour beyond known signatures. More recent AI capabilities tend to support investigation, alert interpretation and parts of containment or remediation, with agentic approaches still developing.
Application security	3	3	3	1	2	LLMs are showing clear value in application security, particularly for reviewing code, identifying insecure patterns, explaining vulnerabilities and suggesting fixes within developer workflows. Agents are beginning to extend this into multi-step testing and remediation, but adoption is still emerging.
Incident Response and Threat Intelligence	2	3	3	2	3	AI agents are increasingly being used to gather context, correlate alerts, summarise incidents and work through investigation steps across multiple tools, while LLMs make threat intelligence and incident data easier for analysts to interpret and act on.
Vulnerability Management	3	3	3	1	3	LLMs are particularly powerful for code-level vulnerability analysis and remediation, helping identify insecure patterns, explain root causes, generate fixes, and validate proposed changes. ML can prioritise vulnerabilities using exploitability, asset context and threat activity, while agents can automate enrichment, ticketing, remediation workflows and follow-up validation.

3 — **Strong fit**: One of the most useful technologies for this use case; proven and deployable today.
2 — **Useful**: Adds meaningful value, usually alongside other approaches rather than as the core technology.
1 — **Emerging / niche**: Relevant in some scenarios, but limited adoption or a relatively narrow use case today.
0 — **Limited fit**: Little practical value for this use case.

*Check [Appendix 4](#) for full taxonomy
** Parker & Lawrence Research analyst assessment. Scores represent practical model fit, not surveyed adoption, product performance or a vendor ranking.

AI Adoption and Risk

AI Deployment: % of firms reporting each level of deployment in Cybersecurity

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
AI is embedded in core operational processes	13%	15%	13%	10%	20%	13%	13%
Multiple production deployments across the domain	28%	25%	33%	32%	27%	13%	30%
Limited production deployment in specific workflows	35%	35%	30%	36%	33%	50%	28%
Pilot or proof-of-concept deployments	16%	15%	23%	17%	13%	10%	18%
Exploring/evaluating use cases	8%	10%	3%	5%	7%	13%	13%
No current AI activity	-	-	-	-	-	-	-

41%
of firms have deployed AI at scale in Cybersecurity

67%
of firms report that AI is causing high levels of risk in Cybersecurity

Cybersecurity leads the taxonomy for AI maturity, with 76% reporting production deployment and 41% operating multiple deployments or embedded AI. That position is no surprise. Cyber has been a “fight fire with fire” market for years: defenders must analyse volumes of network, identity and endpoint activity that exceed human capacity, while attackers continually automate and adapt their methods. High stakes, sustained budgets and a mature vendor market have made AI a standard feature across security products, rather than a separate innovation programme. The headline still requires some caution, however. Only 13% say AI is embedded in core processes, while 35% remain confined to specific workflows, suggesting that buying AI-enabled tools is further advanced than integrating them across the security operating model.

AI Risk: % of firms reporting each level of risk in Cybersecurity

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
5 - Critical AI-related risk 5	11%	12%	8%	15%	7%	3%	15%
4 - Very high AI-related risk 4	26%	23%	40%	25%	20%	17%	33%
3 - High AI-related risk 3	29%	35%	23%	30%	30%	37%	20%
2 - Moderate AI-related risk 2	23%	13%	20%	22%	33%	37%	28%
1 - Low AI-related risk 1	10%	17%	10%	8%	10%	7%	5%
0 - No identifiable AI-related risk 0	-	-	-	-	-	-	-

Cybersecurity is joint-highest for perceived AI risk: 67%* rate it as high or above, including 38%* who consider it very high or critical. AI is accelerating familiar threats by making reconnaissance, social engineering, vulnerability discovery and exploit development faster and more scalable. At the same time, firms are introducing new attack surfaces through foundation models, data pipelines and agents connected to internal systems. Prompt injection is a particular concern because the separation between trusted instructions and untrusted content is not always reliably enforced by the model itself. The NCSC describes it as an architectural weakness of current LLMs, while NIST identifies prompt injection, data poisoning, privacy compromise and model misuse as distinct attack classes. Agentic systems raise the stakes because compromised outputs can lead directly to actions in connected systems.

*Note: Percentages are rounded to whole numbers. Aggregated figures are calculated from unrounded data and may therefore differ from the sum of displayed percentages.

Building the AI stack in risk and compliance

Technology Risk

Technology Risk

Category Deep Dive

	Model Applicability Score					
Subcategory*	Rules-based	Predictive ML	LLMs	Multi-modal	Agents	What this means in practice
Technology Governance	3	2	3	1	2	Rules provide the underlying governance structure through policies, control requirements, approval thresholds and accountability. LLMs can help interpret standards, map requirements to internal controls and generate or review governance documentation, while ML can support risk prioritisation. Agents have potential to coordinate reviews and follow-up, but are less central to the core governance model.
Technology Asset Management	3	3	2	2	2	Rules remain important for classification, ownership and lifecycle controls, while ML can improve asset discovery, deduplication, anomaly detection and risk-based prioritisation. LLMs can help normalise and enrich fragmented asset data, with multimodal AI relevant where evidence includes diagrams, screenshots or configuration outputs. Agentic use is emerging around discovery and remediation workflows.
Third Party Risk Management	3	3	3	2	3	This is a strong fit for several AI types. ML can support vendor risk scoring and continuous monitoring, while LLMs can populate and analyse questionnaires, contracts, policies and assurance reports at scale. Agents can go further by gathering evidence, following up on gaps and coordinating assessments across multiple sources and systems.
Model Risk	3	3	2	2	2	Rules are central to model inventories, approval requirements and validation standards, while ML supports performance monitoring, drift detection and challenger analysis. LLMs are increasingly useful for documentation review, model inventory enrichment and comparing model evidence against policy or regulatory expectations. Agents can support validation workflows, but remain more of an orchestration layer than the core analytical engine.
Resilience	3	3	2	2	3	Rules and scenario frameworks remain important for defining tolerances, dependencies and response requirements, while ML can identify emerging operational risks and analyse large volumes of telemetry. LLMs can help interpret incident, dependency and scenario data, while agents have strong potential for coordinating response, recovery and evidence gathering across complex technology environments.
Technology Testing and Assurance	3	3	3	2	3	AI can broaden and accelerate assurance activity across technology controls. ML can identify anomalies and target higher-risk areas for testing, while LLMs can review control evidence, compare it against requirements and identify potential gaps. Multimodal AI can assess non-text evidence, and agents can increasingly orchestrate testing, evidence collection and remediation workflows.

3 — Strong fit: One of the most useful technologies for this use case; proven and deployable today.

2 — Useful: Adds meaningful value, usually alongside other approaches rather than as the core technology.

1 — Emerging / niche: Relevant in some scenarios, but limited adoption or a relatively narrow use case today.

0 — Limited fit: Little practical value for this use case.

*Check [Appendix 4](#) for full taxonomy

** Parker & Lawrence Research analyst assessment. Scores represent practical model fit, not surveyed adoption, product performance or a vendor ranking.

AI Adoption and Risk

AI Deployment: % of firms reporting each level of deployment in Technology Risk

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
AI is embedded in core operational processes	11%	10%	10%	12%	17%	7%	10%
Multiple production deployments across the domain	28%	35%	25%	24%	37%	30%	25%
Limited production deployment in specific workflows	32%	25%	35%	37%	20%	27%	38%
Pilot or proof-of-concept deployments	18%	20%	15%	19%	13%	20%	18%
Exploring/evaluating use cases	11%	10%	15%	8%	13%	17%	10%
No current AI activity	-	-	-	-	-	-	-

39%

of firms have deployed AI at scale in Technology Risk

60%

of firms report that AI is causing high levels of risk in Technology Risk

Technology Risk sits close to the leading group for AI maturity, with 71% reporting production deployment and 39% operating multiple deployments or embedded AI. Adoption is building across technology governance, model risk and third-party risk, particularly in information-heavy activities such as reviewing control evidence, completing risk assessments, monitoring models and conducting supplier due diligence. These workflows are well suited to generative and agentic AI, which can gather evidence, identify gaps and coordinate follow-up across fragmented systems and teams. Yet only 11% have embedded AI into core processes. Most firms are still improving individual activities rather than connecting inventories, controls, issues and remediation into a continuous view of technology risk.

AI Risk: % of firms reporting each level of risk in Technology Risk

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
5 - Critical AI-related risk 5	7%	2%	8%	7%	13%	7%	8%
4 - Very high AI-related risk 4	19%	17%	13%	24%	10%	27%	15%
3 - High AI-related risk 3	35%	47%	25%	35%	30%	27%	35%
2 - Moderate AI-related risk 2	25%	22%	33%	22%	37%	27%	23%
1 - Low AI-related risk 1	14%	13%	20%	12%	10%	13%	18%
0 - No identifiable AI-related risk 0	1%	-	3%	-	-	-	3%

60% believe AI introduces high or greater material risk in Technology Risk. The concern extends well beyond model performance. AI applications depend on data pipelines, foundation models, cloud infrastructure, APIs and open-source components, often supplied by a small number of providers and with limited transparency for the firms using them. They can also remain hidden inside enterprise software, making it difficult to maintain a complete inventory or understand where decisions are being influenced by AI. Agentic systems add another dimension: once models can access systems and execute tasks, they begin to resemble privileged technology users and must be governed accordingly. Existing model, third-party and technology-risk frameworks remain relevant, but firms will need to connect them far more closely.

*Note: Percentages are rounded to whole numbers. Aggregated figures are calculated from unrounded data and may therefore differ from the sum of displayed percentages.

Riskconnect

Technology Leader: Technology Asset Management



Company profile

Riskconnect's Integrated Risk Platform unifies risk and compliance management across the extended enterprise, enabling organisations to see the full picture of risk, collaborate effectively, and act with confidence. With respect to technology asset management, Riskconnect helps organisations:

- Identify technology assets, dependencies, controls and critical business services
- Assess technology risk against internal risk appetite and regulatory requirements
- Monitor changes across technology assets, controls, incidents and regulations
- Respond through coordinated reassessment, remediation and governance workflows

Headquarters	Georgia, United States
Founded	2007
Founders	Bob Morrell, Antonio Dabraio, Roger Dunkin
Website	riskconnect.com
Industry focus	Agnostic
Regional focus	Global

Customers



Product differentiation

Context-aware materiality assessment

Riskconnect continuously evaluates changes to technology assets, controls, policies, regulations and other connected risk information to determine whether they materially affect the organisation's risk profile and warrant reassessment.

AI embedded into governed workflows

Rather than asking users to determine how AI should be applied, Riskconnect embeds AI directly into technology risk workflows. Deterministic logic, enterprise context and language models work together behind the scenes to support repeatable, auditable risk management processes.

Value drivers*

-  **Faster regulatory interpretation:** More than 80% reduction in time spent understanding unfamiliar regulations.
-  **Lower compliance operating costs:** Lower labor costs by automating manual risk and compliance activities.
-  **Earlier risk detection and resolution:** Fewer risk incidents through earlier identification, continuous assessment, and coordinated remediation.

*Statistics are sourced from vendors, according to their own analysis.

Key product features

Technology risk mapping

Existing risks, controls, causes, consequences and related records are connected automatically to build and maintain a complete technology risk landscape, reducing manual mapping effort.

Policy & regulatory impact analysis

Policies are compared against regulatory requirements to identify gaps, explain why they matter and translate regulatory language into practical actions for technology risk teams.

Continuous materiality assessment

Deterministic logic identifies changes across assets, controls, incidents and regulations. Generative AI evaluates whether those changes exceed the firm's own risk appetite and criticality thresholds, explaining why reassessment is required.

Risk assessment & remediation drafting

Risk descriptions, appetite statements, impact assessments, mitigation plans, playbooks and suggested controls are generated in a consistent format to accelerate technology risk workflows.

AI applications

Model type	Capabilities	Key AI risks / threats mitigated
 Rules-based / symbolic AI Deterministic logic identifies relevant records, detects changes and assembles the context language models reason over. Repeatable, auditable activities remain rules-based.	• Data generation & structuring: Connects risks, controls, causes, consequences and related records into a structured technology risk model. • Scenario evaluation: Identifies changes across connected technology risk information and evaluates whether those changes warrant reassessment against the organisation's own risk appetite and criticality thresholds.	• Lack of traceability: Deterministic processes produce consistent, repeatable and auditable outputs before AI reasoning begins. • AI hallucinations: Language models reason over deterministically prepared enterprise context rather than independently discovering or inferring facts.
 LLMs Retrieval-augmented language models interpret technology risk information using organisational context retrieved from platform data, connected systems and external sources. They support technology risk teams by comparing information, assessing impacts and drafting recommendations within existing governance processes.	• Information retrieval: Retrieves the technology assets, risks, controls, policies, regulations and supporting documentation relevant to a specific technology risk assessment or workflow. • Information summarisation & comparison: Compares organisational policies with regulatory requirements, summarises technology risk information and identifies gaps or changes requiring further assessment. • Structured text generation: Drafts risk descriptions, business impact assessments, mitigation plans, risk appetite statements and suggested controls using retrieved organisational context. • Conversational interaction: Retrieval-grounded assistants answer technology risk process questions using organisational documentation and platform data.	• AI hallucinations: Responses are grounded in the organisation's own technology risk information rather than relying solely on pretrained model knowledge. • AI leaking sensitive data: Prompts and organisational data remain within enterprise controls rather than being exposed to public model training. • Lack of interpretability / transparency: AI-generated recommendations are accompanied by explanations of the reasoning behind them.
 AI Agents AI agents coordinate technology risk workflows by combining deterministic logic, retrieval-grounded language models and enterprise context. They evaluate changes, assess materiality, generate recommendations and orchestrate follow-on actions, while keeping humans responsible for consequential decisions.	• Scenario evaluation: Risk Signal Agent evaluates whether changes across connected technology risk information warrant reassessment against the organisation's own risk appetite, criticality thresholds and governance framework, explaining the reasoning behind every recommendation. • Workflow orchestration: Coordinates multi-step technology risk workflows, moving assessments, remediation activities and policy updates through defined review stages while preserving human approval for consequential actions.	• Lack of interpretability / transparency: Recommendations include the reasoning behind each assessment rather than returning an unexplained decision. • Inadequate human oversight: Agents include defined review points throughout technology risk workflows, with users accepting or declining consequential actions before they are complete.

Neotas

Technology Leader: Third-Party Risk Management



Company profile

Neotas provides third-party risk management and risk intelligence for organisations with complex vendor and partner portfolios. Its platform combines onboarding, questionnaires, configurable risk assessment, screening, continuous monitoring, approvals and remediation in one no-code workflow. The same environment extends into automated and analyst-led enhanced due diligence, bringing sanctions, corporate records, adverse media, court data, OSINT and social signals into source-cited, reviewable cases.

120+ customers, including three of the Big Four

Headquarters	London, United Kingdom
Founded	2017
Founders	Ian Howard, Vipul Mishra
Website	neotas.com
Industry focus	Financial services, private equity, healthcare and regulated industrials, including energy.
Regional focus	Global, with clients across the UK, Europe, North America and the Middle East.

Product differentiation

Risk intelligence inside the lifecycle

Neotas connects multi-source screening directly to onboarding, questionnaires, risk scoring, approvals, monitoring and remediation. Submitted data can create and screen entities, shareholders and directors without re-keying, while high-risk cases can escalate from automated review to analyst-led investigation within the same record.

Source breadth with governed AI

Rather than returning database hit lists, the platform combines premium sources with the open, deep and social web, deconflicts records and preserves citations, match criteria and decision history. Client-defined thresholds, permissions and mandatory human sign-off keep AI recommendations inside a defensible control framework.

Value drivers

-  **Increased productivity or throughput:** Automated screening, deconfliction, summarisation and investigation let teams review larger third-party portfolios without scaling analyst effort linearly.
-  **Reduced manual effort:** Questionnaire data flows into entity creation, screening, risk assessment and reporting, reducing re-keying, duplicate checks and movement between disconnected systems.
-  **Faster onboarding, approvals or reviews:** Low-risk cases can move from intake to a source-cited report quickly, while risk-ranked exceptions are routed to the right reviewer or deeper investigation.
-  **Improved detection of risk, fraud or control failures:** Broader OSINT, social, court and corporate coverage, combined with continuous monitoring, surfaces emerging and connected risks that point-in-time database checks can miss.

Key product features

Integrated onboarding & questionnaires
White-labelled forms, progress tracking and dynamic DDQs collect third-party information once. Responses create entity, shareholder and director records, direct screening and feed the risk assessment and final report.

Multi-source screening & entity resolution
Searches sanctions and watchlists, corporate registries, court records, adverse media, OSINT, social channels and uploaded documents. AI resolves aliases and conflicting records, removes duplicates and presents the match criteria and source evidence for review.

Risk-event triage & continuous monitoring
Related articles can be grouped into a single risk event, reducing duplicate review while retaining the source trail. Live monitoring re-screens the third party across risk domains and triggers alerts, reassessment or remediation as new evidence appears.

Agentic investigation & escalation
For higher-risk cases, specialised agents sequence profile discovery, multi-source research, classification and report drafting. Human validation is built into the workflow, and cases can escalate to Neotas analysts for deeper EDD without leaving the platform.

AI applications

Model type	Capabilities	Key AI risks / threats mitigated
 Rules-based / symbolic AI Client-defined taxonomies, thresholds, permissions and no-code logic govern classification, routing and approval; final risk status remains a named human decision.	• Data generation & structuring: Turns questionnaire responses and source findings into entity, shareholder, director, risk-category and case records within a common schema. • Scenario evaluation: Applies client-configured thresholds, risk frameworks and approval logic to determine match relevance, risk tier, monitoring scope and the next review path.	• Unclear accountability or ownership for AI systems: The client defines the taxonomy, permissions and thresholds, while a named reviewer owns the final risk decision. • Inadequate human oversight of AI-driven decisions: AI recommends and prepares the case, but human validation and sign-off are required before risk status changes.
 Predictive ML Machine learning, NLP and graph analytics power profile discovery, entity resolution, risk classification, deconfliction and event clustering.	• Information retrieval: Discovers likely social profiles, aliases, related parties and records using subject identifiers, network links and source-specific match criteria. • Data generation & structuring: Classifies and de-duplicates posts, articles and records, groups related coverage into events and uses confidence scores to route ambiguous matches for validation.	• Lack of capability or robustness: Configurable thresholds and source-level criteria support controlled deployment; ambiguous cases go to human validation. • Lack of transparency or interpretability: Reviewers can inspect confidence, matching parameters and underlying evidence rather than receiving an unexplained alert.
 AI Agents Specialised, configurable browser agents search for relevant content on social media, distinguishing topics and file types, and download evidence according to the user's instructions.	• Information retrieval: Searches sanctions, PEP, corporate, court, adverse-media, deep-web and social sources in parallel using seed data and resolved aliases. • Scenario evaluation: Confirms the subject, clears false positives and applies the client's framework to classify confirmed findings and assign a red, amber or green status.	• Inadequate human oversight of AI-driven decisions: Humans can monitor the agent's activity live or retrospectively, including watching its every step within a screen-recorded browser. • Lack of traceability from AI outputs to business decisions: Agent-derived findings retain source attribution, decision logs and case-level change history.
 LLMs Model-agnostic language models synthesise case evidence, translate and compare multilingual content, and generate editable, source-cited summaries and reports.	• Information summarisation & comparison: Condenses and reconciles duplicate or conflicting findings, and turns several reports about one risk event into a coherent, reviewable narrative. • Structured text generation: Drafts templated due diligence reports and case summaries in the client's format, retaining citations to the supporting sources. • Information retrieval: Uses case-specific evidence from sanctions, corporate, adverse-media, web, social and uploaded-document sources to ground the generated output.	• AI hallucinations/factual inaccuracy: Every finding carries source attribution so reviewers can test generated claims against the underlying evidence. • Lack of transparency or interpretability: Structured outputs expose sources, match criteria and change history; reports remain editable before export. • AI leaking sensitive data: Access is role-based and the platform is model-agnostic.

Building the AI stack in risk and compliance

Compliance Management

Compliance Management

Category Deep Dive

	Model Applicability Score					
Subcategory*	Rules-based	Predictive ML	LLMs	Multi-modal	Agents	What this means in practice
Horizon Scanning	2	2	3	1	3	LLMs are particularly well suited to turning large volumes of regulatory content into concise, contextualised intelligence. ML can improve filtering, relevance scoring and prioritisation, while agents can continuously monitor sources, investigate developments and trigger downstream reviews or alerts.
Regulatory Obligations Management	3	1	3	1	3	Rules remain important for maintaining structured obligation libraries and deterministic mappings, while LLMs can extract, classify and map obligations from complex regulatory text. Agents can take this further by coordinating updates across obligations, policies, controls and owners as requirements change.
Regulatory Intelligence	2	2	3	1	2	LLMs offer the strongest fit for regulatory research, interpretation, comparison and natural-language querying of trusted sources. Predictive ML can support relevance and materiality scoring, while agents have growing potential for carrying out multi-step research and compiling intelligence around specific regulatory questions.
Regulatory Change Management	3	2	3	1	3	The strongest model is likely to combine deterministic workflows with GenAI and agents. LLMs can summarise changes, assess applicability and identify affected policies or controls; agents can then coordinate impact assessments, assignments, remediation and evidence collection across the change lifecycle.
Compliance Assurance	3	3	3	2	3	This is one of the broadest AI opportunities. ML can identify anomalies, prioritise testing and support risk scoring; LLMs can compare policies, controls and evidence against requirements and identify gaps; multimodal AI can analyse non-text evidence; and agents can orchestrate testing, follow-up and remediation workflows.
Compliance Training	2	2	3	3	2	Generative and multimodal AI can make training substantially more personalised and engaging through tailored content, scenarios, simulations, video and conversational learning. ML can help target training based on roles or risk signals, while agents may increasingly support adaptive learning journeys and automated follow-up.

3 — **Strong fit:** One of the most useful technologies for this use case; proven and deployable today.
2 — **Useful:** Adds meaningful value, usually alongside other approaches rather than as the core technology.
1 — **Emerging / niche:** Relevant in some scenarios, but limited adoption or a relatively narrow use case today.
0 — **Limited fit:** Little practical value for this use case.

*Check [Appendix 4](#) for full taxonomy
** Parker & Lawrence Research analyst assessment. Scores represent practical model fit, not surveyed adoption, product performance or a vendor ranking.

AI Adoption and Risk

AI Deployment: % of firms reporting each level of deployment in Compliance Management

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
AI is embedded in core operational processes	10%	10%	15%	7%	17%	7%	10%
Multiple production deployments across the domain	26%	25%	28%	27%	20%	27%	28%
Limited production deployment in specific workflows	36%	37%	35%	40%	37%	37%	28%
Pilot or proof-of-concept deployments	19%	22%	13%	18%	17%	23%	23%
Exploring/evaluating use cases	9%	7%	10%	8%	10%	7%	13%
No current AI activity	-	-	-	-	-	-	-

36%

of firms have deployed AI at scale in Compliance Management

57%

of firms report that AI is causing high levels of risk in Compliance Management

AI is clearly operational within Compliance Management, but it is not yet reshaping the function end to end. While 72% report production deployment, only 10% say AI is embedded in core processes. Adoption remains concentrated in content-oriented use cases, such as regulatory summarisation and policy drafting, which can be introduced within existing workflows. More workflow-integrated capabilities—including obligation extraction and mapping, translating regulatory interpretation into control change, and connecting execution with assurance—are increasingly available through RegTech platforms, but have yet to be deployed at scale across financial services.

AI Risk: % of firms reporting each level of risk in Compliance Management

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
5 - Critical AI-related risk 5	8%	5%	5%	13%	3%	3%	10%
4 - Very high AI-related risk 4	17%	18%	15%	18%	13%	20%	18%
3 - High AI-related risk 3	32%	28%	25%	31%	47%	30%	35%
2 - Moderate AI-related risk 2	27%	33%	33%	25%	23%	33%	15%
1 - Low AI-related risk 1	13%	13%	20%	12%	13%	7%	13%
0 - No identifiable AI-related risk 0	3%	2%	3%	1%	-	7%	10%

Risk concern is running ahead of deep operational integration. Although AI remains concentrated in specific Compliance Management workflows, 57%* already believe it introduces high or greater material risk. The sensitivity is understandable: an inaccurate summary, missed obligation or weakly evidenced interpretation can propagate into policies and controls, creating compliance exposure well beyond the original AI output. As adoption expands, traceability to source regulation and clear human accountability will be as important as the quality of the model itself.

*Note: Percentages are rounded to whole numbers. Aggregated figures are calculated from unrounded data and may therefore differ from the sum of displayed percentages.

AscentAI

Technology Leader: Regulatory Change Management



Company profile

AscentAI uses artificial intelligence to turn regulatory change into structured, actionable regulatory intelligence. Its platform collects and structures regulatory content, identifies obligations and obligation-level changes, maps them to a firm's regulatory framework, and delivers relevant updates into existing teams, GRC platforms and AI environments.

Headquarters	Chicago, United States
Founded	2011
Founder	Brian Clark
Website	ascentregtech.com
Industry focus	Financial Services: banking, payments and money movement, fintech, lending, digital assets, broker-dealers and wealth management.
Regional focus	United States, United Kingdom, Europe and global regulatory coverage.

Product differentiation

Obligation-level regulatory intelligence

AscentAI treats the individual regulatory obligation as the core unit of change. Its technology extracts requirements from rule sets as discrete objects and compares existing and new obligations to identify what has changed. This creates processed obligation data that can be mapped, assessed and acted on, rather than leaving compliance teams to interpret document-level alerts manually.

Change orchestration in the firm's own operating context

AscentAI connects regulatory change to the customer's operating context and delivers relevant intelligence to the people and systems that need it. The platform identifies which obligations may be affected, maps changes to relevant owners and business units, and routes actionable updates into existing GRC workflows, AI environments and other downstream systems. This allows firms to preserve their preferred operating model while using AscentAI as the regulatory intelligence layer.

Value drivers

 **Operational efficiency:** Obligation-level extraction, comparison and impact assessment reduces the time specialists spend finding amendments, reading rule sets and mapping change.

 **Auditability:** Source-linked obligations, change records, ownership and review history can create a defensible chain from regulatory text to the firm's assessment and downstream action

 **Faster regulatory reporting and evidence production:** Full audit lineage links the regulatory source to obligations, controls, workflows and evidence for faster, more defensible reporting.

Key product features

Automated regulatory data collection and enrichment

Ascent ingests millions of lines of regulatory text and leverages proprietary machine learning models to parse and analyse the data.

Obligation identification

Extracts regulatory requirements from rule sets as individual obligation objects, creating structured data that can be compared and linked to the customer's regulatory framework.

Obligation-level change detection and impact matching

Proprietary comparison engine compares new and existing obligations to detect change, then maps affected obligations to the customer's corporate regulatory map to identify which requirements may need attention.

Workflow delivery through platform, MCP and GRC integrations

Routes relevant change intelligence to responsible individuals and downstream systems. Customers can work inside AscentAI, receive updates in a GRC platform or access trusted regulatory data through an MCP server in their own AI environment.

AI applications

Model type	Capabilities	Key AI risks / threats mitigated
 Predictive ML Purpose-built machine-learning and NLP techniques are used selectively to process and structure high volumes of regulatory information.	• Data generation & structuring: Capable of extracting, normalising and structuring high volumes of regulatory information into machine-readable data. • Information retrieval & classification: Can support the identification, retrieval and classification of relevant regulatory information across large datasets. • Information summarisation & comparison: Can be applied to summarise and compare regulatory information, including identifying semantic similarities, differences and changes. • Scenario evaluation: Can support the evaluation of regulatory information against defined contexts, criteria or scenarios to identify potential relevance or impact.	• Lack of capability or robustness: Models can be applied to bounded regulatory tasks where greater control, consistency and domain specificity are required. Human subject-matter expertise can provide an additional layer of validation and assurance. • AI hallucinations: Grounding AI outputs in captured regulatory source material and structured data can reduce reliance on open-ended generation and improve traceability back to authoritative sources.
 LLMs LLMs are used selectively where natural-language understanding or generation adds value, with outputs grounded in trusted regulatory information and appropriate controls applied for the use case.	• Information summarisation & comparison: LLMs can support the summarisation and comparison of regulatory information, helping users interpret changes, similarities and differences across complex source material. • Information retrieval: LLMs can support natural-language interaction with regulatory information, helping users locate and interpret relevant content across large information sets.	• AI hallucinations: Generative outputs can be grounded in regulatory content collected, structured and enriched by Ascent rather than relying on unrestricted model knowledge or open-web searches.

CUBE

Technology Leader: Horizon Scanning



Company profile

CUBE is an AI-native regulatory intelligence company helping risk and compliance teams track, interpret and act on regulatory change. Built on fifteen years of regulatory expertise and more than ten million datapoints processed annually, CUBE combines AI-native technology with global regulatory data and proprietary network data to enable customers to gather deep intelligence and automate compliance and risk processes.

Headquarters	London, United Kingdom
Founded	2011
Founder	Ben Richmond
Website	cube.global
Industry focus	Financial services, including banking, insurance, asset and investment management and payments.
Regional focus	Global

Product differentiation

Regulatory data foundation

CUBE combines global regulatory content with purpose-built classification models, a regulatory ontology and more than 250 regulatory, risk, research and editorial specialists. The resulting data layer is structured, source-linked and expert-validated, providing an authoritative, current context that models trained on general public data cannot reliably reproduce.

Prioritisation connected to action

Horizon scanning is treated as the first stage of a wider regulatory change process. New developments are evaluated against the organisation's footprint, products, services and regulatory priorities, then can be linked to obligations, controls, impacted teams and downstream workflows without breaking the evidence chain.

Value drivers*

-  **Increased productivity:** One of CUBE's agentic coworker capabilities understands a customer's profile and can accordingly enrich data and prioritise regulatory insights most relevant to them, enabling them to make faster decisions. Customers leveraging the coworker are expected to achieve productivity gains by reducing the time and effort required to triage regulatory updates.
-  **Reduced manual effort:** By applying customer-specific information, CUBE's agentic coworker capabilities reduce the amount of "noise" in horizon scanning by ~60% compared to standard AI-powered solutions, based on early testing.
-  **Improved detection of risk, fraud and control failures:** Global coverage and firm-specific relevance reduce the chance of missing material change and help surface downstream control gaps sooner.
-  **Faster regulatory reporting and evidence production:** Full audit lineage links the regulatory source to obligations, controls, workflows and evidence for faster, more defensible reporting.

Key product features

Global content capture & enrichment

Captures publications from more than 10,000 issuing bodies across 250 countries and territories and 100 languages, then structures, translates and classifies them within a common regulatory schema. Coverage is maintained directly from primary sources, including jurisdictions and languages under-represented in general-purpose training data.

Customer-specific prioritisation

A firm-specific 'passport' records footprint, departments, products and services. This is a persistent, governed representation of the firm that a general-purpose assistant does not retain. A background agent assesses each development for applicability, priority and issuance certainty, surfacing the items most likely to require attention.

Explainable summaries & reasoning

Regulatory-domain language models produce concise overviews and key points, while visible reasoning explains why an item was prioritised. Users can review and amend the output.

Change-to-action traceability

Source links, taxonomy tags, discussion history and workflow records remain connected as a development moves from detection into obligations, controls, remediation and evidence, so the evidence chain, and not only the answer, is preserved for audit.

*Statistics are sourced from vendors, according to their own analysis.

AI applications

Model type	Capabilities	Key AI risks / threats mitigated
 Rules-based / symbolic AI A regulatory ontology and deterministic logic provide governed structure for classification, search, filtering and downstream workflows.	• Data generation & structuring: Applies jurisdiction, issuing-body, business-area and issuance-type tags to standardise incoming content within a canonical schema. • Information Retrieval: Uses the ontology and customer profile to retrieve conceptually relevant content even where regulatory terminology differs.	• AI hallucinations/factual inaccuracy: Higher-level models operate on structured, source-linked regulatory data rather than relying on unbounded web content. • Lack of traceability from AI outputs to business decisions: Sources, tags and downstream workflow records remain attached to each item through review.
 Predictive ML Bounded classifiers handle high-volume tasks where stable, repeatable outputs matter, including issuance-type classification and content labelling. Identical inputs produce consistent, repeatable classifications, a property general generative models do not guarantee.	• Data generation & structuring: Classifies regulatory content into a controlled taxonomy, drawing on tens of thousands of labelled issuance examples. • Scenario evaluation: Identifies the issuance stage, from concept through to in force, to support the certainty dimension used in prioritisation.	• Lack of capability or robustness: Models are trained and evaluated against specialist-curated target data and improved through expert and customer feedback. • AI hallucinations/factual inaccuracy: Bounded classifiers are used for tasks requiring consistent labels rather than open-ended generation.
 LLMs Regulatory-domain language models summarise developments and explain prioritisation using source material and platform context.	• Information summarisation & comparison: Produces concise overviews and key points for new regulatory developments to accelerate review. • Structured text generation: Generates visible reasoning explaining why an item has been prioritised for the organisation.	• AI hallucinations/factual inaccuracy: Outputs are grounded in the source issuance and CUBE's regulatory-domain data rather than the general knowledge held in a horizontal model. • Lack of transparency or interpretability: Users can see the reasoning behind a recommendation, review the source and amend the output.
 AI Agents Background agents combine the firm profile, regulatory content, classifiers and language models to deliver a specific outcome: prioritised horizon-scanning work.	• Scenario evaluation: Assesses applicability and priority against the firm's footprint, departments, products and services, alongside the development's issuance certainty. • Workflow orchestration: Runs as new content arrives, repeatedly triaging the horizon-scanning feed and surfacing risk-ranked work without continuous prompting. • Information retrieval: Combines the regulatory item, taxonomy tags, firm profile and related platform context from multiple sources.	• Inadequate human oversight of AI-driven decisions: The agent recommends and explains; compliance users retain authority over review, escalation and action. • Lack of transparency or interpretability: Each recommendation includes visible reasoning and supporting content. • Lack of traceability from AI outputs to business decisions: Discussion, edits and follow-on workflows remain connected to the original regulatory item.

Gryphon

Technology Leader: Compliance Assurance



Company profile

Gryphon operationalizes contact governance by translating regulatory requirements and enterprise policies into real-time controls at the point of customer engagement. It governs human and AI-driven outreach across channels, determines when and how contact may lawfully proceed, and creates the evidence needed to prove each decision, which allows enterprises to reduce risk without unnecessarily suppressing customer reach.

Headquarters	Boston, MA
Founded	1999
CEO	Clay McNaught
Website	www.gryphon.ai
Industry focus	Financial services, insurance, healthcare, retail and communications; 20+ industries served.
Regional focus	US-first (federal plus all 50 states); Canada, UK and EU for multinational deployments

Product differentiation

The control layer turns regulatory interpretation into deterministic rules

Gryphon translates regulatory requirements into deterministic controls applied at the moment of contact. Consent, DNC, timing, frequency and jurisdiction are evaluated as rules, not predictions, so identical inputs produce the same auditable decision in milliseconds. Those rules are maintained through decades of in-house legal and regulatory expertise and direct engagement with the FTC and state attorneys general, then updated centrally as requirements change so clients do not have to reinterpret and reimplement them independently.

Compliance enforcement that recovers customer reach

Gryphon identifies the channels, exemptions and conditions under which contact may legally proceed, rather than defaulting to broad suppression. Decisions are enforced inside existing CRM, campaign, contact-centre and agentic systems, preventing prohibited outreach while recovering customers who would otherwise be unnecessarily excluded.

Value drivers*

	Revenue recovery: 35–45% of legally reachable contacts recovered from over-suppression, from an audience the business already owns.
	Assurance coverage: 100% of interactions audited and linked to a rule version, replacing sampling-based review.
	Governance Maturity: Benchmarks contact governance maturity and identifies the gaps organizations need to close to advance toward continuous enterprise governance
	Campaign speed: Compliant campaign launch in hours rather than months.
	Enterprise scale: 15bn+ interactions a year governed for a single enterprise client.

*Statistics are sourced from vendors, according to their own analysis.

Key product features

Real-time contact permission decisioning
Evaluates every proposed interaction against regulatory rules, consent, suppression records, timing, frequency, jurisdiction, channel and exemptions. The platform can approve, condition or block contact before engagement.

Omnichannel and agent governance
One rule set applied across voice, SMS, live agents, automated campaigns and autonomous AI agents, with every touchpoint registered and attributed.

Policy-linked auditability
Records the applicable policy, rule version, inputs, requesting user or system and resulting action, enabling each interaction to be reconstructed without relying on generated audit narratives.

Embedded compliance delivery
Makes compliance decisions available through APIs, application integrations and agentic interfaces, allowing controls to operate within CRM, campaign, contact-centre, decisioning and AI-agent workflows.

Continuous governance & control testing
Continuously tests contact-governance controls across campaigns, channels, vendors and workflows while generating evidence for GRC, audit and regulatory response.

AI applications

Model type	Capabilities	Key AI risks / threats mitigated
 Rules-based / deterministic The contact decision, the block, and the evidence record. Explicit, not inferred.	• Information retrieval: Retrieves the inputs a determination requires: applicable regulatory rules, consent records, customer preferences, contact history, suppression status and enterprise policy. • Scenario evaluation: Tests the request against consent, DNC, timing, frequency, jurisdiction, channel, purpose and applicable exemptions, returning permitted, prohibited or conditional. • Workflow orchestration: Returns the determination to connected campaign, CRM, contact-centre and agentic workflows, and blocks non-compliant contact at the network layer. A conditional outcome carries the permitted alternative channel or the approval required.	• AI hallucinations: no generated content enters a contact decision or an audit record, so a plausible-but-wrong output is not a possible failure mode. • Lack of interpretability/transparency: every decision traces to the specific policy and rule version that produced it. • AI exhibiting unfair bias: contact eligibility is rules-based, not scored. • Lack of capability or robustness: the rules engine owns the critical path; AI failures degrade to deterministic behaviour rather than to an incorrect decision. • Unclear accountability: every interaction is attributed to a policy owner, campaign and rule version.
 LLMs Around the path: interpretation, drafting and query. Gated by expert sign-off before anything reaches production.	• Information retrieval: Retrieves relevant material from regulatory, policy, consent and interaction-data repositories to support research and governance queries. • Information summarisation & comparison: Supports regulatory change monitoring by comparing new requirements with existing policies and contact rules to identify areas requiring review. • Structured text generation: Drafts campaign policies, proposed rule language and audit narratives for specialist review. Not applied to the audit record: evidence is captured, not generated. • Conversational interaction: Users can query governance, policy and interaction data. • Hypothesis & scenario generation: Proposes and tests candidate governance scenarios before corresponding rules or campaigns are deployed. These outputs do not control live contact decisions.	• AI hallucinations: Generated analysis is linked to source material and cannot independently create or modify a production contact rule. • Leaking sensitive data: Tenant isolation, data minimisation and redaction are applied before inference. Client data is not used for model training.

Sherlocq

Technology Leader: Regulatory Intelligence



Company profile

Sherlocq is a retrieval-augmented regulatory intelligence platform built on a global regulatory ontology that maps jurisdictions, regulators, licenses, regulated activities and obligations, it answers from the rulebook, not from model memory, across 100+ jurisdictions, refreshed on real time basis.

The platform is organised into three modules: Research, Analyze and Sanctions Intelligence, unified by a shared retrieval and quality-control layer.

Headquarters New York, NY, United States

Founded 2025

Founder & CEO Bhavin Shah

Website sherlocq.com

Industry focus Financial Services, Law Firms, Compliance Consultants, Regulators, Crypto and Digital Assets

Regional focus Global, with primary depth across the US, UK, UAE, Singapore and Hong Kong

Product differentiation

Regulatory research grounded in primary sources

The research module enables cross-jurisdiction regulatory comparisons, mapping requirements across jurisdictions and highlighting governing laws, key requirements and material differences. Sherlocq uses retrieval-augmented generation (RAG) to ground LLM outputs in a curated corpus of regulatory sources, including rules, guidance, enforcement actions, consultations and speeches.

Policy-level gap analysis

The Analyze module enables policy-level gap analysis by comparing uploaded policies against the latest regulatory requirements in a selected jurisdiction. It identifies missing or outdated coverage, highlights where policies may need updating, and provides remediation recommendations, returning each gap against the provision that creates it.

Value drivers*



Operational efficiency: Cross-jurisdiction comparisons that take a team days or weeks are returned in minutes, cutting regulatory research time by up to 70 per cent in Sherlocq's own [pilot / internal] benchmarking. Policy gap assessments are completed in roughly 15 minutes.



Accuracy and performance: Update to 97 to 98 per cent accuracy on a set of [N] expert-reviewed regulatory questions, across the primary financial centres Sherlocq covers.



Regulatory defensibility: Every response cites the rulebook provision it came from, so conclusions can be verified and evidenced rather than taken on trust.

*Statistics are sourced from vendors, according to their own analysis.

Key product features

Cross-jurisdiction regulatory Q&A

Answers questions on requirements across one or several jurisdictions from primary source material, returning a comparison and citations to the governing provisions.

Policy gap analysis

An uploaded policy is compared against the current standard in a chosen jurisdiction and returned as identified gaps with recommended fixes, replacing manual clause-by-clause reconciliation.

Sanctions and watchlist screening

Screens a name against 400 sanctions, watchlist and PEP sources at once via the OpenSanctions API, returning the match with its designation and an exportable, timestamped record.

Access inside existing AI tools

The research corpus is reachable through live connectors into Claude and ChatGPT. Microsoft Copilot and Google Gemini connectors are in development. Analyze and Sanctions Intelligence remain platform-only.

AI applications

Model type	Capabilities	Key AI risks / threats mitigated
 Rules-based / symbolic AI Applies to sanctions screening, which is list-matching against licensed data rather than inference.	• Information retrieval Rules-based fuzzy matching against licensed sanctions, watchlist and PEP data, not a generative process, keeping results deterministic, reproducible and auditable.	• Lack of Transparency or Interpretability: Deterministic matching criteria provide a clear basis for understanding why an individual or entity was returned as a potential match.
 LLMs Retrieval over a curated corpus of regulator-published material, using a third-party frontier model. The corpus, not the model, is the proprietary asset.	• Information retrieval: Locates the applicable law, article and clause across 30+ jurisdictions from Sherlock's curated corpus of primary source material. • Information summarisation & comparison: Synthesises regulatory requirements into structured answers and compares the position across multiple jurisdictions. • Structured text generation: Produces the gap assessment report, comprising prose describing where a policy falls short of the requirement and drafted wording for the fix. • Conversational interaction: Natural-language querying on platform, web, iOS and Android, and through live Claude and ChatGPT connectors.	• AI Hallucinations/Factual Inaccuracy: Generation is confined to the corpus and every answer cites its source provision, so output can be verified at source. • Lack of Transparency or Interpretability: Citations expose the reasoning path; the user can check the provision behind any conclusion. • AI Leaking Sensitive Data: Uploaded material is not used to train models and is subject to defined retention/purge policies. Private-cloud, on-premises and air-gapped deployment options are available where data cannot leave the client environment.

Building the AI stack in risk and compliance

Conduct Risk

Conduct Risk

Category Deep Dive

	Model Applicability Score					
Subcategory*	Rules-based	Predictive ML	LLMs	Multi-modal	Agents	What this means in practice
Trade Surveillance	3	3	2	1	2	Rules and predictive ML remain the core technologies for detecting suspicious trading patterns and known abuse scenarios. ML can identify more complex or previously unseen behaviours, while LLMs can support alert investigation and case summarisation. Agents are emerging as a way to automate evidence gathering, triage and escalation around surveillance alerts.
eComms Surveillance	2	3	3	3	3	This is one of the strongest areas for modern AI. ML and LLMs can identify suspicious language, intent, sentiment and contextual relationships across communications, while multimodal AI extends surveillance across voice, images and other media. Agents can bring together communications, trading activity and contextual evidence to support investigations.
Best Execution	3	3	2	1	2	Rules and quantitative ML are best suited to assessing execution quality across factors such as price, cost, venue and timing. ML can identify anomalous execution outcomes and benchmark performance, while LLMs can help explain findings and support review documentation. Agents may increasingly automate monitoring, investigation and remediation of execution exceptions.
Trade Reporting	3	2	2	1	2	Rules-based systems remain essential for validating and submitting prescribed regulatory data. ML can improve data-quality monitoring and identify anomalous or potentially incorrect submissions, while LLMs can help interpret changing reporting requirements and map them to internal data. Agents can support exception management, remediation and reporting workflows.
Marketing and Communications Compliance	3	2	3	3	3	Generative and multimodal AI have particularly strong potential for reviewing communications before publication. LLMs can assess text against regulatory requirements and internal policies, while multimodal models can review imagery, video and audio alongside written claims. Agents can orchestrate review, approval and amendment workflows.
Customer Outcomes Oversight & Assurance	2	2	3	3	3	Customer outcomes assurance is a hybrid use case: rules set standards and thresholds; ML surfaces risk patterns; LLMs interpret evidence and policy; multimodal AI expands coverage across channels; and agents coordinate testing, monitoring, escalation and remediation.

3 — **Strong fit:** One of the most useful technologies for this use case; proven and deployable today.
2 — **Useful:** Adds meaningful value, usually alongside other approaches rather than as the core technology.
1 — **Emerging / niche:** Relevant in some scenarios, but limited adoption or a relatively narrow use case today.
0 — **Limited fit:** Little practical value for this use case.

*Check [Appendix 4](#) for full taxonomy
** Parker & Lawrence Research analyst assessment. Scores represent practical model fit, not surveyed adoption, product performance or a vendor ranking.

AI Adoption and Risk

AI Deployment: % of firms reporting each level of deployment in Conduct Risk

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
AI is embedded in core operational processes	9%	10%	10%	6%	10%	10%	13%
Multiple production deployments across the domain	18%	13%	13%	15%	27%	13%	33%
Limited production deployment in specific workflows	32%	32%	35%	38%	23%	30%	25%
Pilot or proof-of-concept deployments	28%	30%	30%	28%	23%	37%	18%
Exploring/evaluating use cases	13%	15%	13%	13%	17%	10%	13%
No current AI activity	-	-	-	-	-	-	-

27%
of firms have deployed AI at scale in Conduct Risk

55%
of firms report that AI is causing high levels of risk in Conduct Risk

AI adoption is spread across several largely separate control areas. 59% report production deployment, although only 27% have reached multiple deployments or embedded AI. In trade surveillance, firms are using AI to refine detection and prioritise alerts; in eComms surveillance, LLMs are improving the interpretation of context, intent and coded language; and in marketing compliance, multimodal models and agents are beginning to automate the review and monitoring of customer communications. These are meaningful advances, but they do not form a single integrated workflow in the way that KYC and investigations or regulatory interpretation and obligation management can. Maturity will therefore develop at different speeds across the category.

AI Risk: % of firms reporting each level of risk in Conduct Risk

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
5 - Critical AI-related risk 5	6%	2%	3%	10%	3%	7%	10%
4 - Very high AI-related risk 4	19%	20%	20%	22%	20%	7%	18%
3 - High AI-related risk 3	29%	33%	25%	25%	37%	33%	30%
2 - Moderate AI-related risk 2	27%	28%	38%	28%	20%	30%	18%
1 - Low AI-related risk 1	13%	12%	15%	13%	17%	17%	10%
0 - No identifiable AI-related risk 0	5%	5%	-	2%	3%	7%	15%

55%* believe AI introduces high or greater material risk in Conduct Risk. Within surveillance, firms are introducing non-deterministic models into controls whose decisions must be repeatable and defensible; if AI is prioritising or deprioritising alerts, teams need to understand what it may be filtering out, not only what it surfaces. Beyond the control environment, generative AI allows marketing teams and influencers to produce far more content and variations than traditional approval processes were built to review, while AI-generated rumours, deepfakes and increasingly autonomous trading strategies create new avenues for market manipulation.

*Note: Percentages are rounded to whole numbers. Aggregated figures are calculated from unrounded data and may therefore differ from the sum of displayed percentages.

Aveni

Technology Leader: Customer Outcomes Oversight & Assurance



Company profile

Aveni provides AI-powered oversight and assurance for customer interactions in regulated financial services. Its compliance suite tests customer-facing AI before deployment, monitors live conversations and assesses completed interactions across voice, digital and document channels. Aveni's capabilities are powered by Aveni Intelligence, an in-house developed model layer and SME Insight.



Headquarters	Edinburgh, United Kingdom
Founded	2018
Founders	Joseph Twigg, Professor Lexi Birch, Jamie Hunter
Website	aveni.ai
Industry focus	Financial services, particularly banking, wealth management and financial advice.
Regional focus	United Kingdom (core); international expansion planned.

Product differentiation

Assurance across the AI lifecycle

Aveni evaluates customer-facing AI before, during and after deployment. Approve simulates interactions and tests agents against financial-services risks; Assure monitors every live turn and can allow, escalate or stop a conversation; Detect performs post-interaction QA and evidences customer outcomes across the complete case.

Tiered financial-services model architecture

Aveni orchestrates 150+ models across low-latency classifiers, including Aveni's domain-specific FinLLMs and other frontier models, selecting depth according to the task. Performance benchmarks indicate that the specialist layer delivers a 61% faster time to first token and +3% F1 accuracy versus Gemini Flash 2.5 Lite as well as an overall 92% cheaper running cost (when used in associated use cases).

Value drivers*

-  **Increased productivity or throughput:** In one cited example QA assessment fell from 90 to 15 minutes using Detect, a six-fold increase in review throughput.
-  **Reduced technology costs:** The specialist layer delivered a 61% faster time to first token and +3% F1 accuracy versus Gemini Flash 2.5 Lite as well as an overall 92% cheaper running cost (when used in associated use cases).
-  **Improved detection of risk, fraud or control failures:** Full-population monitoring surfaces vulnerability, conduct, complaint and policy risks that sample-based QA can miss.
-  **Faster regulatory reporting and evidence production:** Each assessment retains reasons and supporting evidence, while Approve produces sign-off-ready packs for governance and regulatory review.

*Statistics are sourced from vendors, according to their own analysis.

Key product features

Pre-deployment agent testing

Connects a customer-facing AI agent via API or SDK, generates realistic personas and edge-case dialogues, and evaluates the resulting conversations against accuracy, safety, suitability and conduct benchmarks before deployment. Results are packaged into sign-off-ready evidence with an audit trail.

Real-time agent guardrails

Monitors each turn of a live customer-facing AI interaction. Fast Tier 1 models detect emerging risks inside the latency budget; escalated cases receive deeper Tier 2 analysis, with configurable actions to allow, alert a human, redact or stop the interaction.

Full-population outcome testing

Assesses complete post-interaction cases across voice, chat, email, documents and related records. Detect applies the firm's QA and Consumer Duty framework to every interaction, surfaces vulnerability, conduct and complaint risks, and preserves evidence for review, reporting and remediation.

Financial-services model intelligence

Combines more than 150 financial-services models across fast classifiers, specialist FinLLMs and selectively used frontier models. Client risk taxonomies, policy frameworks, regulatory knowledge and expert-reviewed evaluation sets govern which model is used, how outputs are tested and where data is deployed.

AI applications

Model type	Capabilities	Key AI risks / threats mitigated
 Rules-based / symbolic AI Client risk taxonomies, policy frameworks and deterministic routing rules define what is tested, how findings are scored and where human intervention is required.	• Data generation & structuring: Encodes regulatory requirements, conduct risks, client policies, personas and thresholds as machine-readable assessment criteria. • Scenario evaluation: Applies configured rules to decide whether an interaction passes, receives deeper analysis, escalates or stops.	• Unclear accountability or ownership for AI systems: The client defines the obligations, thresholds and permitted actions; Aveni operationalises the policy. • Regulatory non-compliance: Assessment criteria remain aligned to the firm's policies and regulatory outcomes rather than generic model judgement.
 Predictive ML Tier 1's fast, task-specific classifiers assess every conversational turn inside the latency budget and surface signals for deeper review.	• Scenario evaluation: Classifies live interaction turns against the client's risk taxonomy, detecting vulnerability, conduct, complaint and policy signals as they emerge. • Data generation & structuring: Labels turns and cases by risk type, severity, outcome and channel to support routing, reporting and trend analysis.	• Lack of capability or robustness: Task-specific models are evaluated against specialist datasets and continuously calibrated through evaluation sets, model cards and expert review. • AI hallucinations/factual inaccuracy: Bounded classifiers handle fast detection tasks where consistent labels matter, avoiding open-ended generation.
 LLMs Tier 2 FinLLMs reason over escalated cases using full interaction evidence and financial-services context; Tier 3 frontier models are used selectively for broader tasks.	• Information retrieval: Retrieves the interaction evidence, policy, regulatory context and prior case information relevant to a specific assessment. • Information summarisation & comparison: Compares the conversation and supporting documents against expected customer outcomes, risk criteria and control requirements. • Structured text generation: Produces evidence-backed reasons, recommended actions and audit-ready assessment narratives for reviewers.	• AI hallucinations/factual inaccuracy: Reasoning is grounded in the complete case and Aveni's financial-services knowledge layer, rather than general model memory alone. • Lack of transparency or interpretability: Tier 2 returns its reasoning, evidence and recommended action for human review. • AI leaking sensitive data: Models can be deployed in the client's private cloud, retaining control of sensitive interaction data.
 AI Agents Approve, Assure and Detect orchestrate multi-step assurance workflows before, during and after deployment, combining simulation, specialist models, evidence and human escalation.	• Data generation & structuring: Approve generates synthetic customer personas and dialogue sets for edge-case testing before launch. • Scenario evaluation: Assure evaluates each live turn and Detect assesses the complete interaction against risk, suitability, conduct and vulnerability criteria. • Workflow orchestration: Routes work from Tier 1 to Tier 2, then allows, flags, redacts, escalates or stops activity according to configured controls.	• Inadequate human oversight of AI-driven decisions: High-risk findings can alert a reviewer or stop the interaction; humans retain authority over consequential decisions. • AI producing harmful content: Real-time guardrails can block, redact or escalate unsafe or non-compliant agent behaviour before the interaction continues. • AI exhibiting unfair bias: Pre-deployment simulations and live monitoring test behaviour across personas and surface emerging bias or unequal outcomes.

Adclear

Technology Leader: Marketing & Communications Compliance



Company profile

Adclear is agentic compliance software for financial marketing. It reviews, clears and monitors marketing and communications for regulated firms, automates regulatory checks against the firm's own rule set, scans live public channels for unapproved changes, and maintains a structured audit trail.

Customers

monzo

NINJATRADER

Pension

coinbase

Headquarters	London, United Kingdom
Founded	2024
Founders	Cameron Ward, Doni Hoti, Joe Jordan
Website	adclear.ai
Industry focus	Financial services: banks, consumer fintech, lending and credit, crypto exchanges, trading platforms, asset managers, insurers
Regional focus	UK and Europe (core), US (active expansion), APAC and rest of world regulatory coverage.

Key product features

Rule and policy library

Regulatory handbooks are broken down into individual obligations and mapped to the firm's own policies, claims and disclaimers. Every rule carries its version history and links back to the source provision.

Pre-publication review

Copy, images, audio and video are checked before release. Separate specialist checks run across the material and return a single comment, each point explained in plain language and tied to the rule behind it. Outputs are audit-ready for regulators such as FINRA, the SEC and NFA.

Conditional self-approval

Firms configure custom approval workflows and the conditions under which a clean automated check is sufficient to release a promotion, by channel, product and jurisdiction. Accountability for the approval stays with the firm's authorised approver, who sets and owns those conditions. Anything outside them routes to the firm's compliance team, and every automatic clearance carries a full audit trail and is verified against the published asset.

Post-publication monitoring

Live material across internal releases, websites, affiliate, social and community channels is matched back to what was approved, and flagged where it has drifted or was never seen.

Product differentiation

Automated throughput

Adclear lets firms set gating rules by channel, product and jurisdiction under which a clean result clears the promotion without human review. The same check is exposed via API, so a client's marketing agents query the rule set as they create and clear content agent-to-agent. Riskier material still surfaces for compliance checks.

Works inside your existing stack

Adclear runs where the work already happens. Reviews are triggered from Figma, Google Docs, Microsoft Word and Adobe Creative Cloud as material is drafted, with findings returned as comments in the file itself. Approval steps can sit inside workflow tools such as Workfront, Monday and Asana, so a promotion moves through the same board it always has, with compliance as a stage rather than a separate queue. Where no native integration exists, the same checks are available through the API and browser extension. Every review, wherever it was started, lands in one record with the same rule references and audit trail.

Value drivers*



Operational efficiency: 88% reduction in approval times. 5-7 days saved per reviewer per month.



Speed to market / throughput: Up to 33x more marketing output.



Content scope: 250,000+ ads reviewed and monitored across pre- and post-publication workflows across 40+ jurisdictions.



First-time approval rates: First-time approval rates typically rise from around 35% to over 90%.

*Statistics are sourced from vendors, according to their own analysis.

AI applications

Model type	Capabilities	Key AI risks / threats mitigated
 Rules-based AI Deterministic logic governs approval routing and gating rather than assessment. Firms configure conditions by channel, product and jurisdiction that determine whether Adclear's clean result permits self-attestation or requires compliance review.	• Scenario evaluation: firm-set conditions on channel, product and jurisdiction test each result to determine whether a promotion clears or escalates.	• Inadequate human oversight: the firm sets where a clean check is sufficient; anything riskier routes to a compliance reviewer. • Unclear accountability: rule authority stays with the firm. Adclear applies configured policy rather than deciding what compliance means.
 AI Agents Task-specific agents perform individual compliance checks against Adclear's Knowledge Hub and other relevant inputs in parallel. Their findings feed an aggregator that produces the single output presented to the reviewer.	• Information retrieval: specialist agents identify the relevant regulatory rules, policies, claims evidence and disclaimers and assess the promotion against them. • Workflow orchestration: separate agents assess disclaimers, imagery, transcript extraction and reading age, with agent activity optionally visible as it runs. • Scenario evaluation: client-side marketing agents query Adclear's knowledge base via API during content creation, receiving feedback before publication. • Information retrieval: agents identify live material, compare it with approved content and surface exceptions for review.	• Lack of transparency: the review is decomposed into named checks rather than presented as a single unexplained verdict. • Lack of capability or robustness: agents are specialised to single tasks and back-tested in-house, so each is deployed against work it has been tested on.
 LLMs Generative AI turns the outputs of Adclear's specialist agents into clear reviewer-facing comments, explanations and suggested amendments.	• Structured text generation: produces the consolidated review comment, explains the issue and suggests corrective wording.	• AI hallucinations: findings are grounded in the underlying checks and linked back to the rule or policy that triggered them.
 Multimodal AI Multimodal review covers the full range of content and media types, from webpages and product flows to imagery, audio and video, pre- and post-publication.	• Media interpretation & summarisation: reads imagery, audio, and video alongside copy, including the words spoken in a video and where they occur.	• Lack of interpretability: findings are explained in plain language beside the asset and link back to the source provision.

Building the AI stack in risk and compliance

ESG

ESG

Category Deep Dive

	Model Applicability Score					
Subcategory*	Rules-based	Predictive ML	LLMs	Multi-modal	Agents	What this means in practice
ESG Data	3	1	3	2	1	Rules remain important for taxonomy mapping, validation and standardisation, while LLMs are particularly useful for extracting and structuring ESG information from reports, filings and other unstructured sources. Multimodal AI can extend this to charts, images and scanned documents. Predictive ML and agents are more peripheral to the core data aggregation and normalisation use case.
ESG Ratings	3	2	2	1	1	Rules remain central to ESG ratings, with defined methodologies determining how indicators are weighted, benchmarked and converted into scores. Predictive ML can strengthen the underlying data through estimation, anomaly detection and signal identification, while LLMs can bring additional context from disclosures, controversies and other unstructured sources. Agents are less central because the core task is still scoring and methodology rather than workflow orchestration.
Portfolio Management	3	3	2	1	2	Rules support exclusion criteria, mandate constraints and portfolio limits, while ML is a strong fit for optimisation, forecasting and risk-return analysis. LLMs can help interpret ESG research and company disclosures, while agents have emerging potential to combine research, screening and portfolio analysis across multiple systems.
ESG Reporting and Disclosures	3	2	3	2	2	LLMs are particularly well suited to drafting, reviewing and mapping disclosures against regulatory requirements and reporting frameworks. Rules remain important for prescribed calculations, templates and validation, while multimodal AI can help extract evidence from varied source documents. Agents can coordinate data collection, review, remediation and approval across the reporting process.
Climate Risk Quantification	2	3	2	2	1	Predictive ML is the strongest fit for modelling climate exposures, scenario outcomes and financial impacts across large datasets. LLMs can support interpretation of scenarios, assumptions and external research, while multimodal AI can contribute where satellite, geospatial or imagery data is used. Agents remain peripheral to the core quantification problem.
Sustainability Management	3	3	2	1	2	Rules provide the structure for targets, methodologies and performance tracking, while predictive ML can model future performance, forecast progress against targets and help prioritise reduction initiatives. LLMs can make sustainability data and plans easier to interrogate and interpret, while agents are beginning to support data preparation, analysis and execution workflows across sustainability programmes.

3 — **Strong fit:** One of the most useful technologies for this use case; proven and deployable today.
2 — **Useful:** Adds meaningful value, usually alongside other approaches rather than as the core technology.
1 — **Emerging / niche:** Relevant in some scenarios, but limited adoption or a relatively narrow use case today.
0 — **Limited fit:** Little practical value for this use case.

*Check [Appendix 4](#) for full taxonomy
** Parker & Lawrence Research analyst assessment. Scores represent practical model fit, not surveyed adoption, product performance or a vendor ranking.

AI Adoption and Risk

AI Deployment: % of firms reporting each level of deployment in ESG

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
AI is embedded in core operational processes	6%	5%	8%	9%	-	-	8%
Multiple production deployments across the domain	13%	18%	15%	14%	10%	10%	8%
Limited production deployment in specific workflows	19%	10%	20%	26%	23%	17%	15%
Pilot or proof-of-concept deployments	33%	35%	45%	27%	23%	43%	35%
Exploring/evaluating use cases	28%	32%	13%	24%	43%	30%	35%
No current AI activity	-	-	-	-	-	-	-

19%
of firms have deployed AI at scale in ESG

32%
of firms report that AI is causing high levels of risk in ESG

ESG remains the least mature category for AI adoption, with 38% reporting production deployment and 61% still at pilot or exploration stage. AI has already demonstrated a strong fit for extracting, classifying and analysing ESG information, but foundational data weaknesses continue to hinder adoption at scale. Firms are still dealing with incomplete disclosures, inconsistent definitions, incompatible reporting formats, estimated emissions and limited data from customers and supply chains. The challenge becomes even greater for Scope 3 emissions, physical-risk exposure and social indicators. AI can make this information easier to process, but it cannot make weak evidence reliable. More advanced applications in reporting, climate analysis and portfolio management are beginning to emerge, but their reliability ultimately depends on getting this data foundation right.

AI Risk: % of firms reporting each level of risk in ESG

	Total	Country					
		UK	France	US	Australia	Singapore	UAE
5 - Critical AI-related risk 5	5%	7%	3%	6%	-	7%	5%
4 - Very high AI-related risk 4	9%	12%	8%	12%	7%	10%	3%
3 - High AI-related risk 3	18%	13%	25%	21%	10%	20%	15%
2 - Moderate AI-related risk 2	30%	35%	28%	25%	27%	23%	43%
1 - Low AI-related risk 1	27%	20%	30%	29%	33%	30%	25%
0 - No identifiable AI-related risk 0	11%	13%	8%	7%	23%	10%	10%

ESG also records the lowest perceived AI risk, with 32%* rating it as high or above and 38%* reporting low or no identifiable risk. That partly reflects the category's earlier maturity and the largely assistive role AI plays today. The principal risk is the loss of the evidential chain: incomplete data can be converted into polished narrative, and modelled estimates into apparently precise ratings, scenarios or disclosures. If the underlying sources and assumptions are unreliable, AI can scale inconsistency and greenwashing as readily as it improves reporting. AI is also becoming an ESG issue in its own right, as firms face growing scrutiny of the energy, water and supply-chain impacts associated with their use of compute.

*Note: Percentages are rounded to whole numbers. Aggregated figures are calculated from unrounded data and may therefore differ from the sum of displayed percentages.



Market direction: The future state of AI in risk and compliance



Trusted autonomy will be built on hybrid AI architectures

Financial institutions have combined deterministic and probabilistic methods for decades. In fraud detection and transaction monitoring, for example, predictive models generate risk scores, while fixed rules, thresholds and workflows determine whether an alert is raised, a transaction is stopped or a case is escalated. The next evolution adds generative intelligence and agentic orchestration to this established pattern.

LLMs add the ability to interpret unstructured information, synthesise context and make bounded judgements. Agents can use that intelligence to coordinate multi-step activity across systems. But the actions available to them – and the conditions under which those actions may be taken – remain governed by deterministic permissions, thresholds, approvals and escalation rules.

In other words, probabilistic models provide flexibility and judgement, and deterministic systems provide consistency, auditability and control. Rather than being displaced by more advanced AI, deterministic infrastructure will be the trust layer that allows it to operate autonomously at scale.

Marc Benioff, Chair and CEO of Salesforce, succinctly captured the logic behind this emerging architecture: **“Probabilistic intelligence alone doesn’t run a company, and deterministic systems don’t reason.”** His example is commercial, but the underlying principle applies. Probabilistic intelligence becomes **more useful** when coupled with deterministic data, rules, permissions and workflows.

“If we can do something with conventional code, that is often the best approach: it is consistent, deterministic and easier to audit. We use that code to gather the relevant context, then feed it into the language model to make the judgement.”

Kim Alderman, Director, AI Innovation Lab, Riskconnect

Human oversight is moving up a level: from reviewing actions to governing autonomy

AI agents are doing more, faster. As they take on more work, “human in the loop” will need to mean something different. If every action requires human approval, the reviewer becomes the bottleneck and many of the benefits of automation are lost. The future is not less human oversight, but a different form of it. Two emerging practices show how this is beginning to work.

First, we are already seeing deterministic decision engines placed between agentic reasoning and execution. Adclear and Gryphon show this model in action, with agents submitting proposed actions to a rules engine that returns a permitted, blocked or conditional outcome to the surrounding workflow. Instead of reviewing every action, humans define the policy that the engine applies repeatedly. This is the deterministic-plus-probabilistic architecture becoming operational infrastructure.

“It is about moving compliance professionals from manual reviewers to compliance engineers: they manage the policy, the policy feeds the AI, and the AI does the heavy manual work.”

Joe Jordan, Co-founder, Adclear

Second, we are beginning to see AI used to supervise other AI. Supervisory systems can monitor activity at machine speed, identify anomalies or policy breaches and direct the most important exceptions to human reviewers. We’re not the only ones recognising this; the **Financial Stability Board** has called out this emerging model as “AI-in-the-loop” and suggests it may become warranted as financial institutions scale their AI use cases. It also identifies “human-on-the-loop,” where people intervene periodically or when required, and “human-in-command,” where they determine the level of autonomy, set guardrails and manage the overall impact.

Human accountability does not disappear under these models. It just moves up the stack, from approving every individual action to setting boundaries, monitoring outcomes, reviewing exceptions and retaining the ability to intervene, which is exactly the model that AscentAI has built. The future loop may therefore contain fewer routine human touches, but more deliberate human control.

“‘In the loop’ suggests that subject-matter experts are touching everything, which can slow the process down. ‘On the loop’ means they step in when the AI reports low confidence and human review is needed.”

Mark Holmes, Head of Solutions Consulting, AscentAI



Outcomes will become the unit of value

AI is reducing the time and cost required to build software features, interfaces and workflows. That does not make software unimportant, but it does make functionality easier to reproduce and less valuable in isolation. Buyers will increasingly judge a product by the completed result: a resolved alert, an approved customer, an evidenced control, a compliant communication or a faster decision.

Agents accelerate this shift because they are designed around goals rather than screens. A conventional application equips a user to perform a task; an agent is expected to combine information, reasoning and action to complete it.

“An agent delivers an outcome. It may use a combination of information, AI and actions to get there, but it is outcome-driven.”

Graeme Wynn, Global Executive Head – Technology, Data and AI, CUBE

In risk and compliance, however, an outcome is not merely an answer. It must be sufficiently accurate for the use case, grounded in the institution's policies and data, traceable to evidence and supported when something goes wrong. Implementation, domain expertise and ongoing assurance therefore become part of the product rather than peripheral services.

Paradoxically, AI may strengthen the role of services. Established vendors' implementation, customer success and subject-matter expert teams can turn general AI capability into repeatable outcomes. AI-native vendors are building similar teams or partnering with consultancies and managed service providers to supply the regulatory understanding that technology alone cannot provide.

“An AI-native start-up can spin up a product quickly; the hard part is having enough domain understanding to challenge a model's very confident answers before the customer does.”

James Mackonochie, Global Executive Head of Product, CUBE

The emerging model is technology-enabled delivery: AI supplies speed and scale, while experts provide context, challenge and rigour. The strongest vendors will not necessarily be those with the most features, but those that can reliably deliver a useful and defensible outcome.

Agents will make end-to-end workflows real

“End-to-end” has long been used to describe platforms that automate several connected steps while still requiring users to move information, switch systems and resolve exceptions manually. Agents can make the term more literal. They can retrieve context, call specialised models, apply rules, update records, trigger actions and assemble evidence across a workflow.

That capability also changes the interface. In many cases, users will not open a separate AI application. The agent will sit inside the case-management, CRM, email or collaboration tools they already use, acting as a digital co-worker across the systems behind them. The visible interface becomes less important than the agent's access, permissions and ability to complete the process.

True end-to-end execution is demanding. Each additional step introduces another dependency, data source, permission and

potential failure mode. The strongest near-term opportunities will therefore be bounded workflows in which routine cases can be completed automatically and exceptions can be escalated cleanly.

“The highest-value agentic workflows are the ones that run routine, low-risk cases end-to-end.”

Ian Howard, Director, Neotas

As agents perform more work, firms will also need a more precise view of cost. Per-seat licence fees reveal little about a workflow that may consume model calls, external data, system actions, monitoring and human review. Cost will need to be traced to the completed outcome and compared with the analyst effort the agent replaces or augments.

The shift is from providing employees with access to software towards governing digital actors that operate across the institution. Orchestration, identity, observability and unit economics will become as important as the intelligence of the underlying model.

Context engines will become core AI infrastructure

The least visible part of an effective AI deployment is often the work that happens before the model is asked to do anything. Platforms must understand the institution's products, jurisdictions, policies, risk appetite, data definitions, organisational roles, thresholds and existing systems. This discovery and onboarding process gives the technology a usable frame of reference.

All enterprise platforms contain some form of context, whether through configurations, prompts, permissions or mapped data. As AI moves from assistance towards judgement and action, that layer becomes more important. A model can be technically capable and still produce the wrong result if it applies the wrong policy, assumes the wrong business model or overlooks the jurisdiction in which a decision is being made.

“A regulatory tool that works for a bank in Toronto may not fit the same bank in London; the business model and regulatory context are different.”

Katsuko Ishizeki-Chaudhari, Expert Advisor
(ex-UK regulator and banking C-suite)


More capable AI does not remove the need for setup. Without reliable context, “more” can mean more output, more corrections, more checking and more cost. AI is not only about automation: getting it aligned with the institution can be anything but automated. The most effective implementations will invest heavily in the unglamorous work of structuring institutional knowledge, resolving ambiguities and defining which sources and rules take precedence.

Over time, this will evolve from one-off configuration into a maintained context engine: a governed, machine-readable representation of the institution that can be reused across models, agents and workflows. The most advanced versions may resemble a digital twin of the control environment, propagating changes in policy, product, jurisdiction or risk appetite across every relevant AI use case.

This context layer will become a major source of both trust and differentiation. Underlying models may become increasingly interchangeable; an accurate, current and institution-specific understanding of how the business actually operates will not.



Market map: AI in RegTech









Appendix

Appendix 1: Survey scope (financial services)

300

final respondents



Base: 300 senior risk & compliance professionals across six markets.



Appendix 2: Survey scope (vendors)

100
vendor respondents

Base: 100 risk and compliance technology vendors.



Appendix 3: Regulatory Prescription and Maturity Methodology

Scope

The scores assess the regulatory environment facing financial institutions that develop, procure or deploy AI, based on legislation, formal regulatory guidance, supervisory initiatives and implementation infrastructure available as at 21 August 2026.

They do not measure the broader strength of a country's AI ecosystem, including investment, computing infrastructure, research capacity or public-sector adoption. Nor do they assess whether a regime is favourable to AI deployment.

The six survey markets are treated as jurisdictions:

- **France** is assessed through the EU AI Act and its implementation and supervision in France.
- **The US** reflects the combined federal and material state-level environment faced by institutions operating across multiple states. It does not imply that every state has the same requirements.
- **The UAE** combines the onshore CBUAE regime with relevant rules and supervisory activity in the DIFC and ADGM. Requirements differ according to an institution's legal domicile.
- **The UK**, Singapore and Australia are assessed primarily at national level, including financial-sector regulation and data-protection requirements.

These are structured Parker & Lawrence Research analyst scores, rather than an external legal benchmark. They should not be interpreted as legal advice.

Regulatory Prescription

Regulatory prescription captures the extent to which firms face direct, legally enforceable and operationally specific AI-related obligations.

The assessment considers:

Component	Weight	Assessment question
Legal force and enforceability	40%	Are the requirements binding, and can non-compliance result in supervisory action, sanctions or civil penalties?
Operational specificity	30%	Does the framework specify controls, documentation, testing, oversight or other procedures that firms must implement?
Breadth and direct applicability	20%	How broadly do the obligations apply across AI systems, firms and stages of the AI lifecycle?
AI-specific supervision and enforcement	10%	Are named authorities, enforcement mechanisms or AI-specific supervisory processes in place?

Score	Regulatory posture
1	Predominantly voluntary principles, with limited direct or enforceable AI-specific obligations
2	Existing-law overlay supported by AI-specific guidance; predominantly principles- or outcomes-based
3	Mixed regime, with binding obligations for certain sectors, states, data types or use cases
4	Broad, binding and risk-based regime with substantial procedural detail, but meaningful flexibility or material scope limitations
5	Comprehensive, binding AI-specific regime with lifecycle obligations, designated supervision and sanctions

A high prescription score does not necessarily mean that a regime is more mature, more effective or less supportive of innovation.

Final scores are expressed on a five-point scale and rounded to the nearest 0.5.

Regulatory Maturity

The Parker & Lawrence Regulatory Maturity Index assesses how far a jurisdiction’s regulatory framework is clear, finance-relevant, institutionally coherent and operationally implementable.

Component	Maximum score	Assessment question
Clarity and completeness	20	Are scope, terminology, timelines and core lifecycle expectations sufficiently developed?
Financial-sector specificity	20	Are financial institutions, regulated decisions and sector-specific risks directly addressed?
Supervisory ownership and coherence	15	Are responsibilities allocated to identifiable authorities, with limited overlap, contradiction or fragmentation?
Implementation infrastructure	20	Are practical guidance, standards, testing facilities, sandboxes, toolkits or assurance mechanisms available?
Supervisory and enforcement readiness	15	Do authorities have the powers, expertise, processes and regulatory coverage needed to supervise the framework in practice?
International alignment	10	Is the framework interoperable with major international standards and regulatory approaches?
Total	100	

Maturity bands

Differences of approximately three points or fewer should not be treated as a meaningful difference in maturity. The US, Australia and UAE therefore sit within the same broad maturity tier, as do Singapore and the UK.

Score	Maturity band
90–100	Very high
75–89	High
60–74	Moderate
40–59	Developing
Below 40	Nascent

Treatment of regulatory evidence

The following rules were applied consistently:

- **Legislation and regulation already in force** receive full weight.
- **Enacted but future-dated obligations** increase prescription, but their staged implementation reduces current supervisory and enforcement readiness.
- **Formal regulatory guidance** contributes to maturity and operational specificity, but only increases prescription in proportion to its legal force.
- **Consultations and proposed rules** indicate direction of travel but do not themselves increase prescription.
- **Toolkits, standards, sandboxes and testing programmes** increase implementation maturity but are not treated as substitutes for binding law.
- **Existing general or sectoral law** is included where an authority has explicitly stated that it applies to AI or automated decision-making.
- A framework is not rewarded simply for containing more documents. The assessment focuses on whether those documents collectively create clear and actionable expectations.



Appendix 4: Regulatory Domains Taxonomy

Financial Crime Compliance

Subcategory	Definition
Identity Verification	Solutions that establish and authenticate an individual's identity using identity documents, biometric checks and liveness detection.
KYC / KYB	Solutions that verify and screen customer (KYC) and business (KYB) identities, including data collection, identity verification, sanctions screening, and adverse media checks.
Fraud Detection & Prevention	Technologies that detect and prevent fraudulent activity by using behavioural analytics, machine learning, and real-time decisioning across customer and transaction data.
Transaction Monitoring	Systems that analyse financial transactions to detect and flag suspicious activity linked to money laundering, terrorist financing, or sanctions evasion.
Sanctions & Screening	Solutions that screen customers, counterparties, and transactions against global sanctions lists and watchlists to identify prohibited or high-risk relationships
Financial Crime Risk Assessment	Solutions that identify, assess, and quantify financial crime risk across customers, products, services, geographies, channels, and business activities, supporting risk-based controls, prioritisation, and ongoing monitoring.

Data Risk

Subcategory	Definition
Data Governance	Solutions supporting the policies, ownership, standards and oversight used to manage data across the organisation, including data catalogues, business glossaries, stewardship workflows, data classification and governance reporting.
Data Lineage and Lifecycle Management	Solutions tracking and visualising the flow, transformation, and lifecycle of data from origin to consumption, ensuring transparency, accountability, and data integrity.
Data Quality	Solutions that monitor, assess, and improve the accuracy, completeness, consistency, timeliness, and validity of data through profiling, validation, cleansing, matching, and issue remediation.
Data Protection	Tools safeguarding organisational data from loss, corruption, or unauthorised access through encryption, backup and recovery systems, access controls, and activity monitoring.
Data Privacy	Solutions managing compliance with data privacy regulations and standards, including consent management, data subject requests, privacy assessments, and breach notification processes.
Data Access and Usage Controls	Solutions controlling who can access data, how it may be used and under what conditions, including entitlement management, policy enforcement, masking, tokenisation and monitoring of data access and usage.

Cybersecurity

Subcategory	Definition
Identity and Access Management	Solutions that control and govern user access to technology systems and data through authentication, authorisation, access provisioning, privileged access controls, and identity lifecycle management.
Network Security	Technologies that protect networks, infrastructure, and data flows from unauthorised access, malicious activity, and cyber threats through monitoring, detection, segmentation, filtering, and access controls.
Endpoint Security	Solutions that protect endpoint devices such as laptops, desktops, servers, and mobile devices by detecting, preventing, and responding to malware, ransomware, unauthorised activity, and other cyber threats.
Application security	Technologies that identify, prevent, and remediate security vulnerabilities across applications and software throughout the development lifecycle, including code analysis, testing, runtime protection, and secure development controls.
Incident Response and Threat Intelligence	Platforms enabling the identification, analysis, containment, and remediation of cybersecurity incidents, enriched by real-time threat data and intelligence sources.
Vulnerability Management	Solutions that continuously identify, assess, prioritise, and track remediation of security vulnerabilities across technology assets, using vulnerability data, asset context, threat intelligence, and risk-based prioritisation.

Technology Risk

Subcategory	Definition
Technology Governance	Solutions that establish and oversee policies, decision rights, controls, and accountability for the use of technology across an organisation, including governance of systems, data, architecture, change, and emerging technologies.
Technology Asset Management	Solutions that identify, inventory, classify, and manage technology assets across their lifecycle, including hardware, software, cloud services, applications, and associated ownership, dependencies, and risk information.
Third Party Risk Management	Solutions that assess, monitor, and manage risks arising from vendors, service providers, cloud providers, and other third parties, including due diligence, control assessment, ongoing monitoring, and remediation tracking.
Model Risk	Solutions that support the governance, validation, monitoring, and control of analytical, statistical, and AI models across their lifecycle, including model inventory, performance testing, documentation, approval, and ongoing risk assessment.
Resilience	Solutions that help organisations identify critical technology dependencies, assess disruption risk, set impact tolerances, test recovery capabilities, and maintain continuity of important services during technology failures or cyber incidents.
Technology Testing and Assurance	Solutions that assess the design and effectiveness of technology controls through testing, evidence collection, control validation, audit support, and continuous assurance across technology and cyber environments.

Compliance Management

Subcategory	Definition
Horizon Scanning	Tools that monitor emerging regulatory, legal, and policy developments to identify potential impacts and support proactive risk and compliance planning.
Regulatory Obligations Management	Solutions that identify, structure, map, and maintain an organisation's regulatory obligations, linking requirements to relevant policies, controls, risks, and accountable owners.
Regulatory Intelligence	Platforms that provide structured insights into regulatory environments, including analysis, interpretation, and contextualisation of regulations to support informed compliance decisions.
Regulatory Change Management	Solutions that help firms systematically assess, implement, and document responses to regulatory changes, involving impact analysis, task assignment, workflow automation, and compliance tracking.
Compliance Assurance	Solutions that evaluate an organisation's adherence to regulatory requirements and internal policies through controls testing, gap analysis, risk scoring, and reporting.
Compliance Training	Solutions that deliver, manage, and track employee compliance education, including role-based learning, assessments, attestations, and reporting to support regulatory and policy adherence.

Conduct Risk

Subcategory	Definition
Trade Surveillance	Tools that monitor orders, trades, and related market activity to detect potential market manipulation, insider dealing, spoofing, layering, and other abusive or suspicious trading behaviours.
eComms Surveillance	Technologies that capture, store, and analyse electronic communications (e.g. email, chat, voice) to identify potential misconduct
Best Execution	Solutions that monitor, assess, and evidence whether client orders are executed on the most favourable terms reasonably available, including analysis of price, cost, speed, likelihood of execution, and venue performance.
Trade Reporting	Technologies that automate the accurate, complete, and timely submission of transaction and trade data to regulators, approved reporting mechanisms, and trade repositories.
Marketing and Communications Compliance	Solutions that review and govern marketing materials, financial promotions, disclosures, and investor communications to ensure they are accurate, fair, appropriately approved, and compliant with regulatory requirements.
Customer Outcomes Oversight & Assurance	Solutions that monitor, test and evidence whether products, services, communications, interactions and decisions deliver fair, suitable and compliant customer outcomes.

ESG

Subcategory	Definition
ESG Data	Solutions aggregating, validating, and standardising environmental, social, and governance (ESG) data from diverse sources, enabling consistent analysis and decision-making.
ESG Ratings	Platforms providing scores or evaluations of companies or financial instruments based on their ESG performance, using transparent methodologies, benchmarking, and industry standards.
Portfolio Management	Tools integrating ESG criteria into investment decision-making and portfolio construction, including screening, risk assessment, scenario analysis, and performance tracking.
ESG Reporting and Disclosures	Technologies supporting the accurate preparation and submission of ESG-related reports and disclosures, aligned with global frameworks and regulatory requirements.
Climate Risk Quantification	Platforms quantifying financial exposure to climate-related risks, including physical and transition risks, through modeling, scenario analysis, stress testing, and financial impact assessments.
Sustainability Management	Solutions that manage and track an organisation's sustainability initiatives, including goal setting, performance monitoring, stakeholder engagement, and continuous improvement activities.

About this research

AI in Risk and Compliance 2026 is an annual research initiative by Parker & Lawrence Research and RegTech Analyst examining how financial institutions are adopting, governing and generating value from AI across risk and compliance.

The research explores both sides of the AI challenge: how firms are using AI to strengthen risk and compliance activities, and how they are managing the new risks, governance requirements and operational challenges that come with its adoption. It brings together perspectives from financial institutions, technology providers and industry experts to provide a clearer view of how AI is being deployed in practice, where it is delivering value, and what is required to scale it responsibly.

About Parker & Lawrence Research

Parker & Lawrence Research is a market research firm focused on AI, risk and compliance. It supports the responsible adoption of emerging technologies by delivering market intelligence that informs the decisions of senior risk and compliance professionals. The team has tracked the RegTech market since 2017, developing proprietary taxonomies, databases and methodologies. During this time, it has worked across the market, supporting global regulators with technology strategy and market intelligence, advising financial institutions on gap analysis and vendor selection, and producing thought leadership with technology vendors.

About RegTech Analyst

RegTech Analyst is the world's leading provider of RegTech information services, B2B media products and industry events. The company informs, connects and advances a network of over 60,000 compliance professionals from financial institutions, technology innovators, investors, and advisory firms. From breaking market intelligence and deep-dive analysis to world-class industry events, RegTech Analyst gives the people shaping the future of financial compliance the information, relationships and analysis they need to move faster and make better decisions.

Acknowledgements

A special thanks to the experts who contributed their individual perspectives and feedback on the themes contained in this report:

**Matt Holmes**

Chair, Responsible AI Working Group,
UKAI Trade Association

**Katsuko Ishizeki-Chaudhari**

Expert Advisor; former UK regulator and
banking C-suite executive

**Ricardo Morais**

Consultant,
SET Advisory

**Reva Schwartz**

Co-Founder, Civitaas; former Research Scientist,
NIST

We would also like to thank the senior risk and compliance professionals who contributed on an anonymous basis.

Contact



Mariyan Dimitrov, Head of Business Operations, FinTech Global
mdimitrov@fintech.global



Michael Lawrence, Co-founder, Parker & Lawrence Research
michael@parkerlawrence.co.uk



Nathan Parker, Co-founder, Parker & Lawrence Research
nathan@parkerlawrence.co.uk