
A decorative graphic consisting of four overlapping circles in orange, blue, and white, located in the middle right section of the page.

2026 GRC Benchmarking Report



How GRC teams are balancing AI adoption
and trust in a complex risk environment

What's Inside the Report



- | | | | |
|----------|-------------------------------|----------|----------------------------------|
| 1 | Abstract & About the Survey | 5 | ROI and the Business Case for AI |
| 2 | The State of AI in GRC | 6 | Third-Party Risk Management |
| 3 | The Manual Work Burden | 7 | What GRC Teams Need Next |
| 4 | Trust, Governance and AI Risk | 8 | Methodology and References |



Abstract & About the Survey



GRC Teams Are Ready for AI, But Trust and Fragmentation Are Slowing Progress

GRC practitioners are under pressure to move faster, increase efficiency and give leadership clearer visibility into risk. AI is beginning to take shape across GRC programs, but most organizations are still testing its potential through pilots or specific use cases, rather than embedding it across workflows. Nearly half of survey respondents said AI is still in the experimental or pilot phase, while only a small share said it is embedded across workflows.

The report finds that GRC practitioners are embracing AI, but proving value tends to be uneven across organizations.

The report's core finding is that GRC practitioners are not resisting AI. Instead, they see clear opportunities to reduce repetitive administrative work, improve throughput and spend more time on understanding critical insights. However, the barriers are trust and time. Data privacy, inaccurate outputs and uneven ROI continue to influence how quickly organizations are willing to expand AI use across risk and compliance functions. When combined with the extensive allotted time required of GRC professionals to make this technology functional, AI implementation often remains a low priority.

About the Survey

The findings also point to a broader operational challenge: fragmented work. Respondents identified documentation, risk assessments, evidence collection and third-party reviews as some of the most time-consuming activities in their GRC programs. Those tasks become harder when data, workflows and ownership are scattered across disconnected systems. As AI adoption grows, organizations will need more than new tools. They will need trusted processes, connected data and clear governance structures that allow GRC practitioners to scale AI responsibly.

About the Survey

This report is based on responses from 126 GRC practitioners and related professionals about AI adoption, manual work, governance maturity, third-party risk management and the business case for AI-enabled GRC.

The survey was conducted from March 4 to April 10, 2026, and was designed to benchmark how organizations are using AI today, where inefficiencies still slow GRC programs and what practitioners need to move from experimentation to trusted adoption.

Respondents represented a range of functions, including information security/IT, risk and compliance, finance/audit, operations/procurement, legal/general counsel and other related GRC roles. Survey participants came from organizations across software and technology, finance and banking, manufacturing, business services, education, healthcare, retail, insurance and energy and utilities. The respondent base was primarily located in North America.



The State of AI in GRC

Adoption is moving forward, but
maturity remains uneven.




AI Adoption in GRC

AI Use Is Taking Shape, But Most Programs Are Still Early

AI has entered the GRC conversation, but most organizations have not yet reached mature adoption. Among respondents, 44.4% said AI is still in the experimental or pilot phase and 25.4% said it is used for specific tasks. Another 16.7% said AI is not in use at all, while only 13.5% said it is embedded across workflows.

AI Adoption Remains Largely Experimental



These findings suggest that many GRC functions are beyond initial awareness, but have not yet reached broad operational use. AI is being evaluated, tested and applied in targeted areas, but most organizations have not fully connected it to the workflows that define day-to-day GRC execution.

Early Stage Adoption

This early-stage adoption pattern reflects the nature of GRC work. Risk, compliance and audit processes depend on accuracy, documentation, defensibility and oversight. A cautious approach is understandable, especially when AI outputs may influence risk assessments, evidence review, control monitoring or executive reporting.

The opportunity now is to move from isolated experimentation to practical adoption. For GRC practitioners, the question is less about whether AI has value and more about where it can be applied safely, consistently and measurably. The strongest path forward will likely start with specific, repeatable workflows where AI can reduce administrative work while maintaining human review and governance.

Similarly, PYMNTS Intelligence, a research platform that reports on digital transformations and payments news, found that risk and compliance remain in the early stages of next-generation AI adoption, with use largely confined to exploratory efforts or limited pilots in areas where readiness barriers are more visible. As organizations navigate how to scale AI throughout GRC, they may benefit from evaluating how mixed operating environments impact adoption.



GRC Priorities and Compliance Complexity

GRC Teams Are Operating Across a Mixed Compliance Environment

Respondents identified SOC 2 as the highest-priority industry framework, indicating a possible emphasis on information services. However, the broader spread across frameworks suggests GRC practitioners are operating in a mixed compliance environment. Most organizations are not managing one clean set of requirements. They are balancing multiple frameworks, stakeholder expectations and both internal and regulatory reporting demands all at once.


That complexity amplifies the pain/overhead of manual work processes, as there is no economy of scope or scale.

Each framework may require its own documentation, evidence, testing cadence and reporting structure. When those activities are managed across disconnected workflows, GRC practitioners can lose significant time gathering information, confirming ownership and reconciling data across systems.

These findings lay an important foundation: GRC functions operate within complex requirements, and that complexity becomes harder to manage when repetitive work depends on fragmented systems and manual follow-up.

The Manual Work Burden

The clearest pain point is the time GRC teams lose to repeatable work.



Time-Consuming GRC Activities



Manual Work Is Still Consuming the GRC Function

The survey findings show that GRC practitioners are still spending a significant amount of time on necessary, but highly repetitive, manual work. Respondents identified evidence collection and documentation as the most time-consuming activity, representing 25.9% of selections. Risk assessments followed at 19.8%, while third-party reviews accounted for 16.5%.

These are not minor administrative tasks. Evidence collection, documentation, assessments and vendor reviews are central to how organizations prove compliance, evaluate exposure and respond to audits. The issue is that these activities often require practitioners to gather information from multiple sources, follow up with stakeholders, validate records and prepare materials for review.

That work creates a drag on the GRC function. Time spent collecting, documenting, reviewing and chasing information is time that cannot be spent analyzing risk, improving controls or advising leadership on an organization's security posture. As requirements increase, manual processes make it harder for GRC practitioners to keep pace without adding more burden.

The data points to a clear need for operational relief. GRC practitioners need workflows that reduce unnecessary repetition, make information easier to access and create more time for analysis and related insights. In practice, GRC teams need centralized systems that map frameworks and maintain reusable evidence repositories. This would allow practitioners to reuse evidence without repeatedly requesting the same information from business stakeholders whenever audits, reviews or assessments occur. As a result, governance teams could improve visibility across organizational workflows. This could establish the foundation for automation that simplifies repetitive follow-up processes.

Microsoft's State of the SOC research offers a useful parallel from the broader security field, noting that many SOC analysts spend significant time manually aggregating data across tools. While SOC and GRC functions are distinct, the comparison reinforces how fragmented, manual processes can slow security and risk operations.



Where AI Can Reduce Manual Work



Respondents See AI's Biggest Value in Repeatable GRC Operations

Respondents were clear about where they see the greatest opportunity for AI: reducing repeatable administrative work.

Nearly 70% said AI would have the greatest impact on simplifying repeatable, administrative GRC operations.

That finding separates practical AI value from AI hype. GRC practitioners are not looking for AI to replace professional judgment. They are looking for support in the areas where repetitive work routinely slows audits, reviews, assessments and reporting. These are the tasks that require consistency, documentation and follow-through, but not always deep analysis at every step.

AI's near-term value in GRC may come from helping practitioners perform these repetitive tasks more efficiently. That could include summarizing information, identifying missing documentation, helping prioritize follow-up, surfacing patterns or improving the flow of routine tasks. The common thread is time. If AI can reduce the administrative lift around GRC work, practitioners can spend more time interpreting findings, assessing risk and supporting better decisions.

The Strongest Use Cases Will Still Require Human Oversight

GRC work depends on context, judgment and accountability. AI can help reduce friction, but practitioners still need to validate outputs, confirm accuracy and ensure that results are defensible. For organizations moving from pilot programs to broader adoption, that balance between automation and review will be critical.

External research validates these findings. ISACA notes that AI can help streamline administrative processes and allow employees to focus on more strategic work, but organizations still need protections to ensure output accuracy. Establishing safeguards often starts with identifying where legacy processes are holding efficiency back. In many organizations, inefficiencies are often the result of fragmented GRC workflows.

Fragmentation and Administrative Drag



Fragmented Workflows Are Holding Back GRC Efficiency

The survey findings point to a broader challenge beneath the surface of day-to-day GRC work: fragmentation. Respondents identified evidence collection, documentation, risk assessments and third-party reviews as the most time-consuming activities in their programs. These responsibilities depend on information, workflows and ownership spread across multiple systems or stakeholders, making already complex work more difficult to complete efficiently.

Fragmentation shows up in the everyday mechanics of GRC work. A practitioner may need to collect evidence from one system, validate ownership in another, track follow-up in a spreadsheet and report findings through a separate process. Each handoff creates room for delays, errors, duplication or incomplete information. Over time, those small points of friction add up.

Why This Matters


GRC programs are expected to do more than support compliance. They are required to fortify an organization's entire security posture. Leadership is looking for clearer visibility into risk, faster reviews and stronger decision support. Those outcomes are difficult to achieve when practitioners spend too much time searching for information or reconciling data across systems.

The findings support a broader need for connected GRC operations. Reducing administrative drag requires better alignment across data, workflows and accountability. Without that connection, AI adoption may improve isolated activities but fail to address the root causes of inefficiency.

External industry reports reflect this reality. PYMNTS Intelligence found that only 15% of senior technology leaders describe their data as mostly integrated with few silos, while 65% say their data is moderately integrated and 20% acknowledge fragmentation across teams and systems. The same report notes that fragmented data systems suggest many organizations are not yet equipped to scale AI across the enterprise.

Trust, Governance and AI Risk

GRC teams need AI they can
defend, not just AI they can access.



The Trust Barrier



Privacy and Accuracy Are Limiting AI Adoption

Trust is one of the clearest barriers to expanding AI use in GRC. Respondents identified data privacy and accuracy/hallucinations as the top concerns limiting AI adoption, at 28.6% and 25.4%, respectively.

These concerns reflect the realities of GRC work. Practitioners are responsible for information that must be accurate, documented and defensible. A flawed AI output can create opportunities for compliance risk, misstate exposure, weaken audit readiness or cause leadership to make decisions based on incomplete information.


Because GRC workflows require the exchange of sensitive internal data, third-party information, control evidence, policies and risk findings, organizations need confidence that AI-enabled processes will protect that information and support appropriate access controls.

Accuracy concerns are equally important. GRC practitioners need to know when AI outputs can be trusted, when they require review and how those outputs were generated. Without that confidence, AI adoption may remain limited to low-risk tasks or isolated pilots.

The findings suggest that GRC practitioners are ready to embrace AI, but only if it can be governed, reviewed and defended. The next phase of adoption will depend on whether organizations can build trust into the way AI is introduced, monitored and scaled.

These insights are reinforced by broader industry notions. ISACA's AI risk guidance notes that as AI adoption increases, organizations must account for security vulnerabilities, misuse and a rapidly expanding governance and compliance environment. ISACA also states that risk management now requires continuous oversight and governance across the AI lifecycle.

AI Governance Maturity



Governance maturity, however, does not automatically mean AI is ready to scale. Policies can define acceptable use, review requirements and ownership, but GRC practitioners still need confidence in the data, outputs and workflows behind AI-enabled processes. A governance document will not solve issues tied to fragmented information, unclear process ownership or inconsistent review.

The strongest AI governance programs will connect policy with execution. That means defining how AI can be used, who is accountable for outputs, how results are validated and where human oversight is required. GRC teams should outline operational controls, such as how practitioners will identify and inventory AI usage, monitor and manage data input to AI tools and identify appropriate usage across models. Part of this practice should include understanding which AI use cases are appropriate now and which require greater maturity before broader adoption. Governance leaders must then work to enforce these policies organization-wide.

The survey findings show meaningful progress, but they also reinforce a central point that trusted AI adoption depends on more than interest or policy. It requires governance, connected workflows and operational discipline working together.

Organizations looking to start this process should examine external frameworks. For example, ISACA frames effective AI governance as the integration of risk management into AI design, deployment, monitoring and lifecycle controls, with the goal of preserving trust and accountability.

AI Governance Is Moving from Planning to Practice

AI governance is becoming a priority for GRC practitioners. More than one-third of respondents, 35.7%, said their organization has clearly defined governance for AI use within GRC. Another 34.1% said governance is in progress. Smaller shares said they have informal guidance (18.3%) or no governance yet, but plan to create it (11.9%).

Together, these findings suggest that many organizations recognize the need for structure. AI is no longer being treated only as an emerging technology discussion. It is becoming part of formal governance planning, especially in functions where accuracy, accountability and compliance obligations matter.

ROI and the Business Case for AI

AI interest is high, but measurable return is still catching up



Onspring


AI ROI Is Still Uneven



Many GRC Teams Have Not Yet Proven AI's Return

Many organizations are investing or experimenting with AI in GRC, but proof of value remains uneven. Only 16.7% of respondents report seeing demonstrable ROI from AI investments in their GRC programs, while the majority (43.7%) say they have not yet seen returns. An additional 15.1% of respondents are unsure whether they have received value from the tool, and 24.6% say they have received limited yields.

While these findings point to a disconnect between AI value and investment, the real gap lies in how GRC practitioners measure ROI. Governance teams continue to struggle to define AI value metrics that extend beyond traditional cost-reduction efforts, but ROI is increasingly determined by operational improvements. The organizations that see value are more likely to measure it through speed, throughput and better decisions than through immediate cost savings.



AI ROI isn't absent. It's underdeveloped. If organizations invested their AI in the everyday workflows that slow GRC practitioners, they would see measurable improvements in business efficiency and coordinated processes, while reducing the number of repetitive tasks that drag on operations.

However, realizing measurable AI value extends beyond adoption. It depends on defensibility. Organizations will need clearer governance frameworks, secure guardrails, centralized data repositories and connected operations to begin the process of building trust in AI outputs. Once this occurs, AI can effectively bridge operational gaps. In time, reduced inefficiencies and improved data coordination will likely create a downstream impact on ROI.

How GRC Teams Measure AI Value

Near-Term ROI Is Operational First

Among respondents who reported AI ROI in their GRC workflows, most of that value first appeared through operational indicators. Of those, 25% of respondents cited reduced business cycle time, 20.7% experienced higher throughput of GRC tasks and 19% reported improved decision-making.

This data aligns with broader report findings. Organizations that experience AI value measure it in how the tool handles repetitive, rule-based processes. If AI can give GRC professionals time to focus on higher-value responsibilities that require human analysis and input, they will likely see downstream gains in consistency and efficiency across governance practices.

Over time, these operational gains will likely translate to faster audits, quicker reviews, higher task completion and clearer decisions. These are all tangible productivity boosts that GRC leaders can measure. For many organizations, these workflow improvements may represent where AI investments will make the greatest impact. The remaining challenge is operationalizing it.

Most limitations on realizing AI ROI stem from internal preparation, as noted by external research. A PYMNTS Intelligence report found that 71% of enterprise leaders acknowledge their own readiness is the binding constraint for AI, suggesting that AI value depends on whether organizations have the systems, workflows and data structure to support it.

What Would Justify More AI Investment

Faster Audits, Clearer Risk Visibility and Less Documentation Burden Lead the Business Case

GRC leaders can only justify investment in AI when it delivers clear operational value. 28.8% of respondents reported wanting faster audits and reviews, 24.8% wanted clearer visibility into top risks and 22.4% cited a desire for more time spent on analysis rather than documentation.

These findings point to how AI-enabled GRC investment can create practical value within day-to-day governance processes.

Practitioners want speed, visibility and time back for higher-value work. This suggests that organizations would benefit from AI implementations that prioritize workflow throughput and surface relevant insights for decision-making. That means using AI to reduce documentation burden and improve task flow.

Desired Outcomes of AI Adoption in GRC

Faster audits and reviews

28.8 %

Clearer visibility into top risks

24.8 %

More time spent on analysis

22.4 %

The Future

By embedding AI into governance functions, AI investment could focus on work quality and decision support, assisting GRC practitioners with their operational demands. This would eventually increase productivity, reduce redundant administrative workload and create opportunities for practitioners to focus on judgment-based responsibilities that require human oversight.

Third-Party Risk Management

Confidence is moderate, but visibility and follow-through remain difficult.

Confidence in Third-Party Risk Management



Third-Party Risk Confidence Is Present, But Not Strong


The survey findings indicated that a majority (52%) of GRC respondents felt somewhat confident in their third-party risk management. However, only 25.6% felt very confident in this category, while 10.4% felt unsure and 12% did not feel confident at all. This indicates room for improvement in external-facing GRC domains that could encounter significant risks.

While respondents aren't overwhelmingly negative about their third-party risk management, the data does not suggest strong confidence either. This implies that most organizations have a working process, but still struggle to maintain visibility over time.

External reports echo this reality. Deloitte insights indicate that this lack of oversight frequently comes from targeted, siloed risk management. Usually, organizations assess individual third-party risks by business function. This narrow view removes broader oversight of business exposure, which is critical for understanding overall third-party risk touchpoints.


This limited governance model may become harder to sustain amid evolving threats. For example, recent IBM reports indicate that supply chain and third-party risks have increased sharply over the past five years, placing greater pressure on GRC teams to maintain comprehensive assessments of external and interconnected risks. This trend will likely continue, leaving organizations exposed if they do not adapt their third-party risk management. Limited management of business exposure could have severe operational implications, especially in fields that handle sensitive information.

Mitigating third-party risk will likely depend on organizations building stronger visibility into ongoing exposure. Organizations may benefit from an operational shift toward coordinated and continuous oversight of external touchpoints.



Third-Party Risk Gaps

Continuous Monitoring Is the Top Third-Party Risk Challenge



The GRC respondents identified their most prevalent third-party risk issues: 28.6% cited ongoing/continuous monitoring, 21.4% completion of vendor questionnaires, 15.9% reporting and follow-up, 10.3% accountability and 10.3% evidence review.

These findings suggest that third-party risks are difficult to monitor because GRC teams must maintain oversight of external systems they cannot control.

With increased AI usage among third-party partners, practitioners also face growing concerns about how outside organizations maintain data privacy and compliance standards. Bringing monitoring, follow-up, evidence review and accountability together is challenging in this environment, as GRC teams have to evaluate these external signals to proactively understand risk exposure.

Remedying this issue may require GRC leaders to establish clear monitoring practices across third-party activities and define clear surveillance handoffs between internal partners. These improvements may enhance monitoring, follow-up and accountability across the third-party risk lifecycle.

This sentiment extends beyond this report's findings, with Ernst & Young's (EY) 2025 Global Third-Party Risk Management Survey pointing to similar themes. A growing trend among EY's survey respondents was centralizing third-party risk management to reduce friction points. 52% of these respondents found that this centralization increased understanding of risks during the decision-making process, suggesting that establishing clear visibility practices is a useful method to reduce third-party risk issues. EY's findings also suggest that AI could further accelerate this movement toward centralized third-party risk management.




AI's Role in Managing Third-Party Risk

AI Has Not Yet Transformed Third-Party Risk Management

Only 15.1% of GRC respondents found that AI has transformed their third-party risk management. 31% reported that AI was in pilot for this function, 16.7% stated that it has not transformed this area and 37.7% felt unsure about AI's role in this category or that the question was not applicable to their AI usage.

The limiting factor in AI adoption for this function is less about interest and more about operational complexity. Organizations will likely need stronger governance policies, established guardrails and connected information before GRC leaders can trust this tool to accurately assess and surface insights on third-party risk findings.

AI Transformation in Third-Party Risk Management




Without connected systems and centralized visibility, organizations may struggle to apply AI consistently across reviews, assessments and evidence collection/documentation. This may reinforce many of the existing operational bottlenecks in third-party risk management, including inventory maintenance, survey conduct, response review and overall risk monitoring.

These findings coincide with broader themes throughout this report. AI has the potential to affect third-party GRC relationships, but trust, workflow integration and measurable value must mature before its impact becomes widespread.



What GRC Teams Need Next

The next phase of AI-enabled GRC depends on trust, connection, and measurable outcomes.




Leadership Evaluation of GRC Effectiveness

Leadership Is Measuring GRC Through Efficiency, Audit Outcomes and Risk Reduction

Leadership evaluates GRC effectiveness through several key factors: 26.7% of respondents cite operational efficiency, 24.4% audit outcomes, 24.4% risk reduction, 13.4% informal judgment and 11.1% undefined methods.

How Leadership Measures GRC Effectiveness



The findings suggest that organizations may benefit from outlining leadership success metrics

much earlier in the development of GRC programs. Many governance initiatives start by implementing a process and then defining what value looks like to leadership. This creates a disconnect between what leadership expects and what practitioners prioritize.

To achieve results, GRC leaders may consider ways to reduce friction in typical governance workflows. This means meeting practitioners where they are. Right now, GRC operations revolve around manual documentation and disconnected follow-up, which consumes time and detracts from higher-level analysis and workflows. Focusing on improvements that give practitioners time back and reduce manual burden will allow workflow outcomes to align more closely with what leadership measures.

Actionable Insights

From AI Experimentation to Trusted GRC Execution

The survey findings demonstrate how AI usage in GRC workflows depends on more than just interest, but on addressing the operational challenges hindering execution. Outlined below are several clear operational priorities for GRC teams as they work toward AI implementation.

Core Takeaways:

- 1 Reduce manual burden where it's most visible.**
Start with repeatable GRC operations that consume the most time, including documentation, risk assessments and third-party reviews.
- 2 Build trust before expanding AI across workflows.**
Privacy, accuracy and audit defensibility should shape how AI is introduced, reviewed and scaled.
- 3 Connect systems before expecting stronger ROI.**
AI ROI will be limited if data, tasks and ownership remain fragmented.
- 4 Treat third-party risk as an ongoing process.**
Continuous monitoring, follow-up and accountability need stronger operational support, not just point-in-time questionnaires.
- 5 Measure AI value through practical GRC outcomes.**
Reduced cycle times, higher throughput and improved decision-making may be the strongest early indicators of AI value.

Behind the Data



Methodology

This report is based on survey responses from GRC practitioners collected between March and April of 2026. The survey was designed to benchmark how organizations are using AI within GRC programs, where manual work continues to create operational burden and what factors are shaping the move from experimentation to trusted adoption. Respondents were asked about AI use, governance maturity, manual work, ROI, third-party risk management and leadership evaluation of GRC effectiveness. Supporting research from external sources (included below) and Onspring internal data informed the context and analysis presented in this report.

Behind the Data

Margin/Limitations:

Findings are based on self-reported survey responses and should be interpreted as directional benchmarks rather than a statistically representative view of the entire GRC market. Some questions allowed respondents to select multiple answers, so totals may not equal 100%. Percentages may also vary slightly due to rounding. If the survey was distributed primarily through Onspring's network or customer and community channels, the findings may reflect the perspectives of that audience more strongly than the broader market.

Sample Details:

- Respondent count: 126 respondents.
- Fielding dates: March–April 2026.
- Respondent roles: GRC practitioners, including IT, risk and compliance, finance, operations and legal counsel functions.
- Seniority level: Manager, VP or C-suite roles, as well as director-level roles.
- Industries represented: Highest concentrations in software/technology and finance, with limited responses across manufacturing, business services, education, healthcare, retail, insurance and energy/utilities.
- Geography: Highest concentrations were in North America, with limited responses from other geographic areas.
- Company size: Highest respondent concentrations were in companies with over 5,000 employees and between 201 and 1,000 employees, while other company sizes saw more limited representation.

External Sources:

- [ISACA](#). "The Promise and Peril of the AI Revolution: Managing Risk." March 25, 2026.
- [PYMNTS Intelligence](#). "The Enterprise AI Readiness Gap: What Company Data Reveals About the Real Barrier to Scale." 2026.
- [EY](#). "2025 EY Global Third-Party Risk Management Survey." 2025.
- [KPMG](#). "Global AI Pulse: Q1 2026." April 15, 2026.
- [Microsoft Security](#). "State of the SOC. Unify Now or Pay Later: What New Research Reveals." 2026.
- [IBM](#). "X-Force Threat Intelligence Index 2026." February 25, 2026.
- [Deloitte](#). "Risk Angles: Third-Party Risk." 2015.

