BENCHMARK REPORT

The 2026 GRC Benchmark Report

An annual industry view — framework adoption, audit prep, vendor risk, AI governance, tooling, and program economics.

May 2026 · Edition v1 — Public-data synthesis · Published by Talarity

Contents

About this report

This is the first edition of the Talarity GRC Benchmark. It exists because the current public benchmark landscape is fragmented: a vendor survey here, an industry-association survey there, an analyst press release on AI somewhere else. No single document gives a working GRC leader a citeable, cross-cut view of where programs actually sit in 2026.

This edition is a **synthesis report**. Every quantitative claim is footnoted to a public source. Where the data simply does not exist in the public record, we say so, instead of inventing a number. A subsequent edition (planned for H2 2026) will layer in primary survey data fielded by Talarity directly with GRC practitioners; that data is not yet in this version, and nothing here is presented as if it were.

If you find an error, an outdated stat, or a stronger source for a claim we make, write to **research@talarity.com**. We will correct it in the next edition and credit the contributor.

Executive summary

Eight findings shape the 2026 GRC landscape:

1. **Multi-framework is now the default, not the aspiration.** 97% of organizations conduct at least two compliance audits per year, and 74% of large enterprises run four or more.[1]
2. **The framework portfolio is consolidating around a stable core.** ISO 27001 adoption has reached 81% among organizations engaged with continuous-compliance platforms, and ISO 27001:2022 is now mandatory for all new certifications.[2]
3. **Audit prep remains stubbornly expensive, but automation is finally landing.** Roughly half of practitioners using compliance-automation platforms report cutting prep time by 25–50%, and another 36% report cuts greater than 50%.[2]
4. **Third-party risk is the single weakest link.** 97% of organizations experienced at least one supply-chain breach in 2025 — a 20% year-over-year increase — while 63% of TPRM teams describe themselves as understaffed.[3]
5. **The vendor questionnaire is a fiction the industry keeps signing off on.** 81% of respondents say three-quarters of their vendors pass security questionnaires; only 14% trust that the answers reflect reality.[3]
6. **AI policy lags AI use by a wide margin.** Only 37% of organizations have an AI governance policy, while 80%+ of employees already use unapproved AI tools and IBM measures a $670,000 premium on breaches involving shadow AI.[4][5]
7. **The compliance team is the cost center.** Across surveyed GRC budgets, 32% goes to headcount, 20% to recruitment, and 19% to tooling — 71% of the budget before a single auditor is paid.[6]

8. **GRC tooling is centralizing, but not yet consolidating.** 86% of organizations now run GRC through a centralized team, and 56% use a common controls framework — but the underlying tool stack is still fragmented across audit, vendor, incident, and policy systems.[7]

The thread running through all eight: **programs are doing more compliance with the same staff, on more frameworks, under more scrutiny, against a vendor base that is itself an attack surface, while AI is simultaneously the productivity unlock and the next-largest unmanaged risk.**

Methodology

This v1 edition draws on three classes of evidence:

Recurring industry benchmark studies with disclosed methodology — primarily A-LIGN's 2026 Compliance Benchmark Report (n=1,043 global respondents, fielded Aug–Sep 2025),[1] Hyperproof's 2026 IT Risk and Compliance Benchmark Report (sixth annual edition),[7] ISACA's State of Cybersecurity 2025 (n≈4,000),[8] the SCCE Cross Industry Compliance Staff and Budget Benchmarking Survey,[9] and Cycore's Annual GRC Budget Survey 2025.[6]

Analyst and standards research — Gartner's published research on GRC tools and AI governance platforms,[10] the NIST AI Risk Management Framework, ISO/IEC 42001:2023, the EU AI Act, ISO/IEC 27001:2022 transition requirements, and EU NIS2 enforcement data.

Vendor research with disclosed methodology — Secureframe's 2026 trends data,[2] IBM's 2025 Cost of a Data Breach (shadow-AI cost premium),[4] Acuvity's 2025 State of AI Security,[5] Komprise's 2025 IT Survey,[5] and audit-cost data triangulated across multiple vendor pricing guides where independent sources converge on similar ranges.[11][12][13]

What is **not** in this edition: Talarity-fielded primary survey data, and Talarity-anonymized program telemetry. Both are in scope for the next edition. We have flagged sections where primary data would strengthen the analysis — "what we cannot yet measure" — rather than filling those gaps with estimates.

Citation policy: every percentage, dollar amount, and time figure in this report is footnoted to a source URL. Ranges quoted as consensus across multiple vendor guides are footnoted to the strongest individual source and noted as triangulated.

Section 1 — Framework distribution

Headline findings

The number of frameworks per program is rising, and the underlying portfolio is more standardized than the marketing of any single framework suggests.

A-LIGN's 2026 study finds that **97% of organizations now conduct at least two audits annually, and 74% of large enterprises conduct four or more**.[1] That is the single most important framing statistic in this report: when practitioners talk about "audit fatigue," the median program is no longer prepping for one audit a year. It is prepping for two to four, on rolling cycles, with overlapping evidence demands.

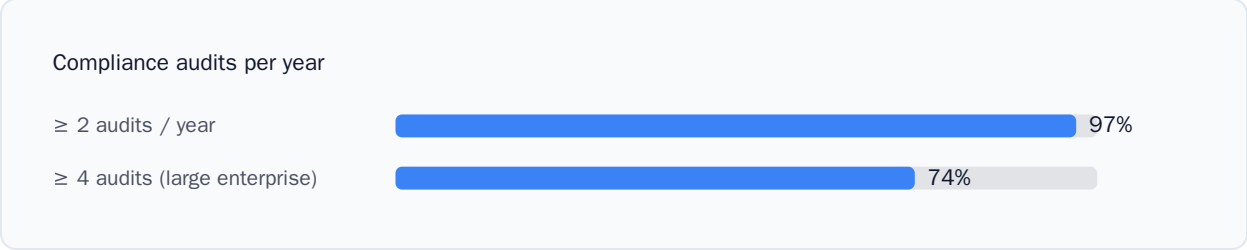

Figure 1 — 97% of organizations now conduct at least two compliance audits per year; 74% of large enterprises run four or more. Source: A-LIGN 2026 Compliance Benchmark Report.

ISO 27001 has crossed the threshold from "international option" to "default international standard." Secureframe's 2026 trends report places ISO 27001 adoption at 81% among organizations using continuous-compliance platforms,[2] and the underlying market is being valued in industry research at $18.59 billion in 2025 with growth projected to $74.56 billion by 2035.[2] Two regulatory forces are amplifying this: ISO 27001:2022 is now mandatory for all new certifications (the 2013 version is no longer accepted), and the EU NIS2 directive is driving ISO 27001 demand across European supply chains as a baseline for regulated entities and their vendors.[2]

The federal compliance stack has gone mainstream. A-LIGN reports that **94% of surveyed organizations are pursuing compliance with one or more of CMMC, FISMA, FedRAMP, or GovRAMP.**[1] That is a remarkable figure — federal frameworks have historically been a defense-contractor specialization. The 94% number reflects three years of CMMC ramp pressure and FedRAMP's broadening reach into commercial SaaS that sells to federal agencies.

The geographic split is real and persistent

The SOC 2 / ISO 27001 split correlates strongly with buyer geography. US and Canadian enterprise procurement still requests SOC 2 first; EU, UK, and APAC procurement requests ISO 27001 first.[2] The interesting evolution: US enterprise buyers are increasingly accepting ISO 27001 in lieu of, or alongside, SOC 2 — particularly for vendors with European footprints — though SOC 2 remains the default first request in most US B2B procurement processes.[2]

For programs running both: 65–75% of controls overlap between SOC 2 and ISO 27001, which is the structural reason common-controls-framework adoption (covered in Section 5) is climbing.[2]

What this means for program design

If you are a GRC leader building a five-year framework strategy in 2026, the data argues for three structural decisions:

1. **Plan for multi-framework from day one.** A single-framework program is now a transitional state. Architecture choices — control catalog, evidence model, audit cycle — should assume two-to-four concurrent frameworks within twenty-four months.
2. **Anchor on ISO 27001 if you have any international ambition.** The combination of NIS2 enforcement and the ISO 27001:2022 transition has made it the most leverageable single certification for cross-border programs.
3. **Treat the federal stack as plausible, not exotic.** With 94% of surveyed organizations now engaged with at least one federal framework, the historical “we’ll get to CMMC if we ever sell to the government” stance is increasingly out of step with where mid-market and enterprise programs are landing.

What we cannot yet measure (deferred to next edition)

- Distribution of frameworks **by company size and revenue band**, beyond the rough enterprise vs. SMB split available in public data.
- Year-over-year framework growth rates at the individual-framework level (SOC 2 → ISO 27001 → HIPAA → PCI DSS → NIST CSF). Public research reports adoption at a point in time; isolating clean YoY growth requires primary data.
- The “framework debt” question: how many programs hold certifications they are no longer actively maintaining at full rigor.

Section 2 — Audit prep timelines

Headline findings

Audit preparation remains a quarter-to-half-year project for most frameworks, and platform automation is the single largest variable in program cost.

For **ISO 27001**, public guidance converges on roughly **150 hours of internal time plus 80 hours of external assistance** for a mid-sized company through the gap-assessment and documentation phase.[14] Small organizations with dedicated support typically complete the gap assessment and documentation in two to four months, followed by Stage 1 and Stage 2 audits in another two to three months. Large enterprises with distributed scope routinely run six to twelve months or longer before the audit window opens.[14]

For **SOC 2**, the public consensus is **two to twelve-plus months total**, with Type 1 typically completing in two to four months and Type 2 reports running six to twelve months because of the longer observation window. Prep itself — scoping, gap closure, control documentation, evidence collection setup — typically runs four to twelve weeks before the observation window even begins.[11] [15]

Automation is finally moving the curve

This is where the data has changed materially in the past two years. Secureframe's 2026 customer analysis reports that **nearly half of its customers reduce audit preparation time by 25 to 50%, and another 36% prepare for audits in less than half the time.**[2] More striking: organizations using AI-driven control mapping and automated evidence collection have reduced **audit expenditures by up to 35%**.[2]

Independent corroboration: Hyperproof's 2026 study finds **76% of GRC professionals still spend 30% or more of their time on repetitive administrative work**[7] — meaning automation has improved the picture for a meaningful minority of programs but has not yet displaced the modal experience of "compliance is mostly evidence-chasing."

The multi-audit multiplier

The cost story changes dramatically once a program is running concurrent audits — which, per Section 1, is now the median state. A-LIGN's data shows **80% of respondents rating the quality of a compliance report as "extremely important," up from 70% in 2025**, and **60% say they would change auditors to improve report quality** — with **83% reporting clear differences in quality**

between audit providers (up from 72% in 2025).[1] When a single program produces four audit reports a year, quality differences compound: a report that an enterprise buyer treats as suspect costs the seller real revenue.

The breach correlation

Hyperproof's 2026 data contains one of the cleanest correlations in the report: **among organizations that experienced a security breach, 58% anticipate spending more time on IT risk and compliance in 2026, compared with only 37% of those who did not experience a breach.**[7] More striking is the direction of causation implied by their other finding: **50% of respondents managing risk on an ad-hoc basis experienced a breach in 2025, compared with 27% among those with an integrated, automated approach.**[7] Mature programs aren't just spending less time on compliance — they're getting breached less often.

What we cannot yet measure (deferred to next edition)

- Median audit prep hours **by framework count**, isolated cleanly. Public data tracks per-framework prep; the question of "what does adding a second or third audit actually cost in marginal hours" requires primary data.
- Hours spent on **evidence collection vs. control design vs. auditor liaison**, broken out. This breakdown is the single most-requested benchmarking question we receive from practitioners and is the strongest argument for primary survey data.
- The cost of **failed first audits** — the implicit tax on programs that ship to audit too early.

Section 3 — Vendor risk practices

Headline findings

Third-party risk management is simultaneously the most-acknowledged weakness in GRC and the least-resourced.

The 2025 numbers are stark. **97% of organizations experienced at least one supply-chain breach in 2025 — a 20% increase from 2024.**[3] **63% of TPRM teams report being understaffed.**[3] **80% of organizations plan to hire more TPRM staff in 2025.**[3] And — the single most damning operational statistic — **94% of organizations are not assessing all the vendors they would like to assess, citing resource constraints.**[3]

The questionnaire credibility crisis

The most uncomfortable finding in any 2025 TPRM study: **81% of respondents say three-quarters of their vendors pass security questionnaires, but only 14% of respondents trust that third-party security actually matches the responses given.**[3] In other words, the dominant mechanism for vendor security validation — the security questionnaire — is something that 86% of TPRM professionals quietly believe is not telling them the truth.

This is not a marginal trust gap. It is a structural indictment of an industry-standard practice. The reasons are not mysterious: questionnaires are filled out by vendor sales engineers under deal pressure, against a control framework the asker rarely audits, and reviewed by a TPRM team that 63% of the time is too understaffed to verify anything.

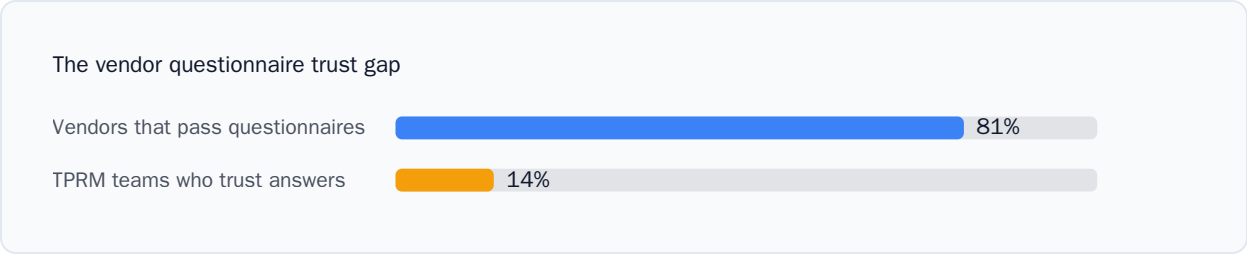

Figure 2 — 81% of vendors pass security questionnaires; only 14% of TPRM teams trust the answers. Source: Atlas Systems — 120+ Third-Party Risk Management Statistics.

The shift to continuous assessment

The mature-program response is well-documented in Hyperproof's research: a shift away from "annual questionnaire + annual reassessment" toward **continuous monitoring combined with structured reassessment cadence**.[7] The pattern Hyperproof describes — dynamic review cadences that adjust frequency based on onboarding stage, performance history, and incident patterns — is the operational mirror of what mature financial-risk programs adopted a decade ago: stop treating risk as point-in-time, start treating it as ongoing telemetry.

The barrier is not philosophical. It is tooling and headcount. Continuous monitoring requires a vendor inventory with metadata richer than most programs maintain, a risk scoring model that updates on signal, and a workflow that escalates to a human only when material change is detected. Programs that hit all three are the minority; programs that achieve any one of them often label themselves "mature" because they have moved past pure questionnaire-driven practice.

The AI-supplier dimension

The newest pressure on TPRM is the AI subsupplier. As vendor ecosystems expand and AI-enabled suppliers become more common, the existing questionnaire model — already not trusted — is being asked to certify behavior of subsystems whose training data, model provenance, and downstream data handling are themselves opaque.[7] Public research has not yet produced clean adoption data on AI-specific vendor questionnaires, but the regulatory pressure is clear: enterprise procurement increasingly lists ISO/IEC 42001 in due-diligence questionnaires, particularly in financial services, healthcare, and public sector.[5]

What this means for program design

Three operational recommendations the data supports:

1. **Tier ruthlessly.** If 94% of programs cannot assess every vendor they want to, the only honest response is explicit tiering with documented criteria, not aspirational coverage.
2. **Reduce reliance on the questionnaire as the primary signal.** Pair it with at least one independent signal — external attack-surface monitoring, breach feed monitoring, evidence-based control review — for every tier above "low."
3. **Add AI-specific questions to the questionnaire now, even imperfectly.** Programs that adopt early — even with rough ISO 42001-aligned questions — will be ahead of the regulatory curve when EU AI Act high-risk obligations become enforceable on August 2, 2026.[16]

What we cannot yet measure (deferred to next edition)

- The actual **distribution of assessment cadence in practice** — annual vs. event-driven vs. continuous — by vendor tier. Public research tells us continuous monitoring is increasing; it does not yet tell us what percentage of vendors at each tier are actually monitored continuously.
- The **gap between vendor policy and vendor evidence** at scale — a question that requires direct program-data access.
- TPRM **cost per vendor assessed**, by program size.

Section 4 — AI in GRC

This is the section of this report where the gap between practice and policy is widest, the regulatory landscape is moving fastest, and the next twelve months will produce the most change. We treat it in three parts: AI in the GRC workflow (the workflow itself), AI governance for internal use (the policy side), and AI governance for vendor AI (the procurement side).

Part 1 — AI inside the GRC workflow

GRC has become one of the most-AI-saturated functions inside the enterprise. **Hyperproof's 2026 study reports 97% of respondents use AI to streamline their work.**[7] **ISACA's 2025 cybersecurity workforce study reports that 47% of practitioners have helped develop AI governance (up from 35% the prior year), and 40% have been involved in AI implementation (up from 29%).**[8]

Demand for AI governance and model-risk skills rose **81% year over year in 2025**.[5] That figure is one of the cleanest signals in this report about where GRC career investment should land.

The interesting nuance from Hyperproof: **the benefit accrues most when AI is embedded directly inside the platform that owns controls, evidence, and assessments — not when AI is bolted on as a disconnected helper.**[7] This matches the structural argument for tooling consolidation in Section 5.

Part 2 — Governance of internal AI use

Here is the gap that defines 2026. The numbers tell a coherent and uncomfortable story.

Only 37% of organizations have AI governance policies, and only 37% have policies specifically to manage AI or detect shadow AI, per IBM's 2025 Cost of a Data Breach Report.[4] Meanwhile, **over 80% of employees already use unapproved AI tools**, and security researchers have tracked **665 distinct generative AI applications** across enterprise environments.[5] An independent survey range puts the proportion of employees using non-IT-approved AI between **40% and 65%**.[5]

A-LIGN's 2026 study corroborates the policy gap from the other side: **33% of surveyed organizations don't have an AI compliance strategy in place at all.**[1]

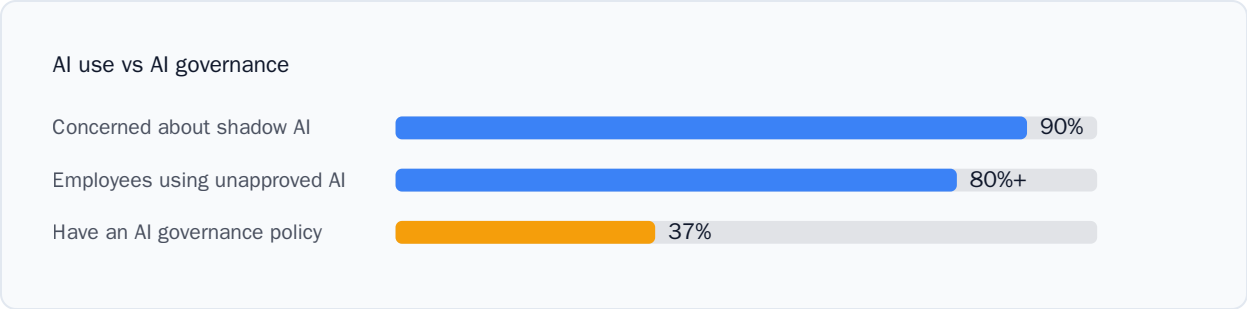

Figure 3 — Enterprises are broadly concerned and broadly using AI; only a third have governance policies. Sources: Komprise 2025 IT Survey, Acuvity 2025 State of AI Security, and IBM 2025 Cost of a Data Breach Report.

The financial impact is now measurable. **IBM's 2025 Cost of a Data Breach Report finds that shadow AI adds $670,000 to the average breach cost**, and insider risk driven by AI negligence has been measured at **$10.3 million annually**.[4][5]

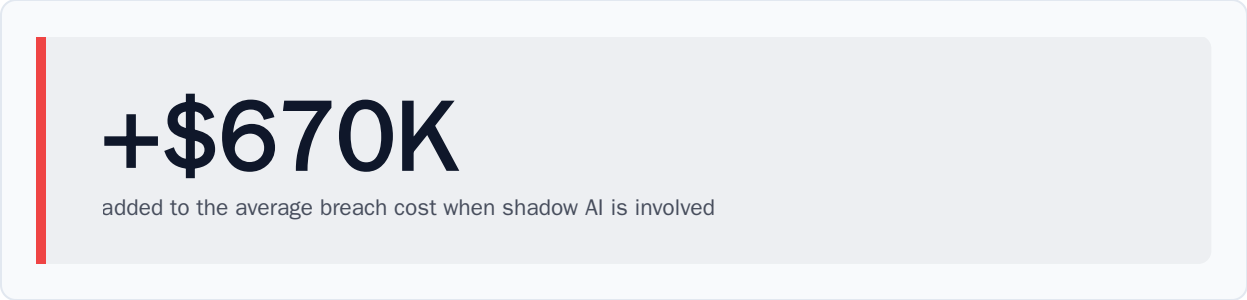

Figure 4 — The measured premium shadow AI adds to a breach. Source: IBM 2025 Cost of a Data Breach Report.

Control implementation is uneven and lags policy. Per public research aggregated in early 2026:

- **57%** of organizations have an acceptable use policy for AI tools
- **55%** have access controls for AI agents and models
- **55%** have AI activity logging and auditing
- **48%** have identity governance for AI entities [5]

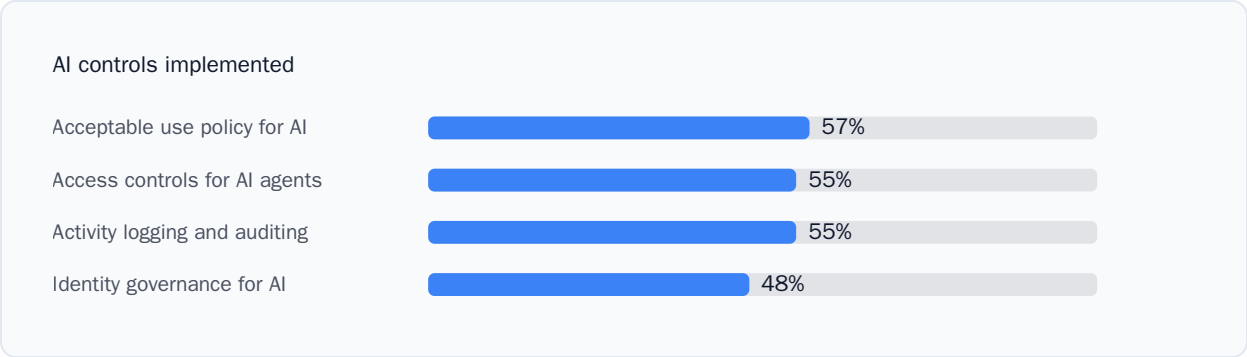

Figure 5 — AI control implementation lags awareness across every category. Source: MarkTechPost — Enterprise AI Governance in 2026, aggregating industry research.

90% of organizations report being concerned about shadow AI from a privacy and security standpoint per Komprise's 2025 IT survey,[5] and Gartner's November 2025 analysis (n=302 cybersecurity leaders) projects that **by 2030, more than 40% of enterprises will experience security or compliance incidents linked to unauthorized shadow AI.**[10]

The takeaway: enterprises are **broadly aware**, **moderately policied**, and **lightly controlled** on internal AI use. The gap between concern (90%) and control (48–57%) is the operational center of gravity for AI governance investment in 2026.

Part 3 — Governance of vendor AI

The procurement side is moving on a clearer regulatory timeline.

The EU AI Act's high-risk obligations become enforceable on August 2, 2026.[16] That is the single date that should be on every GRC leader's wall.

ISO/IEC 42001 is increasingly listed in enterprise procurement due-diligence questionnaires, with adoption pressure concentrated in financial services, healthcare, and public-sector buyers. Public research describes it as **shifting from a differentiator to a procurement requirement**.[5]

The framework landscape itself is converging rather than diverging. Practitioner research consistently observes that NIST AI RMF, ISO/IEC 42001, and the EU AI Act can be satisfied with a single coordinated set of processes, policies, and documentation — with NIST AI RMF serving as the US technical reference, ISO/IEC 42001 as the global compliance layer, and the EU AI Act as the enforceable legal regime for the EU market.[5][16]

What "good" looks like in 2026

Synthesizing across the public data, a program that is credibly governing AI use in 2026 has the following posture:

1. **A written AI acceptable use policy** in production (joins the 57% who have this).
2. **Inventory of AI use** across the organization — both sanctioned and shadow — refreshed at least quarterly. Public research has no clean number on inventory completeness; this is one of the highest-value primary-survey questions for the next edition.
3. **Access controls and identity governance** for AI services (joins the 48–55% range).
4. **Activity logging** for sanctioned AI use (joins the 55%).
5. **Vendor AI clauses** in the standard MSA / DPA — including model provenance, training data attestations, sub-processor notification, and right-to-audit language tied to ISO/IEC 42001 or NIST AI RMF.
6. **A documented EU AI Act readiness plan** with August 2, 2026 as the milestone if the organization sells to or operates in the EU.

The depressing complement of this list is that **no public research suggests most programs hit even three of these six** today. AI governance is the section of GRC where the policy and tooling industry is most actively forming — and where the gap between "we have a policy" and "we have a controlled program" is widest.

What we cannot yet measure (deferred to next edition)

- Distribution of AI governance maturity **by industry**, beyond the rough financial-services / healthcare / public-sector lead.
- The proportion of programs that have **completed an EU AI Act gap assessment** in advance of the August 2026 enforcement date.
- The actual rate of **AI-specific clauses in standard vendor contracts** — a question that almost no public research has touched.

Section 5 — Tooling consolidation

Headline findings

GRC is centralizing organizationally faster than it is consolidating technologically.

Hyperproof's 2026 study finds that **86% of organizations now run GRC through a centralized team**, with only 14% managing GRC through individual teams or business units.[7] And **56% of organizations now use a common controls framework (CCF)** to streamline GRC processes, a meaningful jump that reflects the same multi-framework reality described in Section 1.[7] Common controls frameworks are the structural prerequisite for evidence reuse across SOC 2, ISO 27001, HIPAA, and the other portfolio frameworks that the median program is now running.

The market structure

Gartner observes that the GRC tooling market currently contains **more than one hundred vendors**, and that the analyst expectation is for material consolidation as buyer requirements clarify.[17] The Gartner Magic Quadrant for GRC Tools, Assurance Leaders, and the corresponding Market Guide, remain the canonical analyst references for this segment.

Independent industry analysis published in late 2025 observes that **enterprises now manage on the order of a dozen disconnected GRC applications** — one for audits, one for vendors, one for incidents, one for policy management — and that legal and compliance departments are expected to **increase GRC tooling investment by 50% by 2026.**[18]

The unified-vs-fragmented cost case

Quantitative evidence on consolidation savings is scarce, and the most-cited published case study comes from a single vendor analysis: **Censinet documents an enterprise moving from $2.8 million in fragmented spending to a $750,000 unified platform — roughly $2.05 million in annual savings.**[19] We cite this as a single illustrative case rather than a benchmark; the lack of broader public data on consolidation cost deltas is one of the largest gaps in the existing benchmark literature, and one of the most important primary-survey questions for the next edition of this report.

What is well-documented is the **administrative tax** of the current fragmented model: **76% of GRC professionals spend 30% or more of their time on repetitive administrative work.**[7] If a fragmented tool stack is the structural cause of that admin tax — and there is strong practitioner consensus that it is — then consolidation's return is principally measured in time recovered, not subscription cost saved.

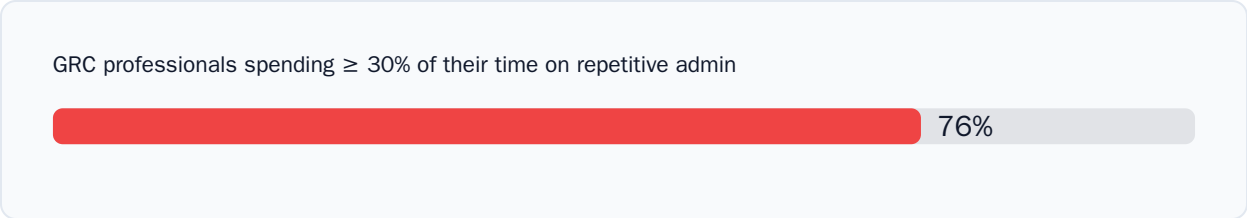

Figure 6 — The administrative tax that defines the modal GRC workday. Source: Hyperproof 2026 IT Risk and Compliance Benchmark Report.

What this means for program design

Three observations:

1. **Centralize the team before consolidating the tools.** The 86% centralization figure suggests this sequencing is already the dominant pattern. Tool consolidation works in a centralized team; it tends to fail in distributed-ownership programs.
2. **A common controls framework is the leverage point.** The 56% CCF adoption number is rising, and it is the prerequisite for evidence reuse across the multi-framework portfolio that is now the median program shape.
3. **Beware of consolidation theater.** Replacing five point tools with one suite that is itself five sub-products with separate data models is a rename, not a consolidation. The Hyperproof finding that AI value accrues most when embedded in the platform that owns controls, evidence, and assessments[7] applies symmetrically to consolidation: the platform that genuinely unifies the data model is the platform that delivers the time savings.

What we cannot yet measure (deferred to next edition)

- The **actual median GRC tool count per program**, by company size. Anecdotal “dozen disconnected apps” figures are widely repeated; a real distribution does not exist in public research.
- **Net savings from consolidation** across a representative sample, beyond the Censinet case.
- **Failure rate of consolidation projects** — particularly the proportion of “consolidations” that produce a more-expensive multi-product stack inside one vendor logo.

Section 6 — Program economics

Headline findings

Headcount, not tooling, is the dominant cost of running a GRC program.

Cycore's Annual GRC Budget Survey 2025 finds that **32% of GRC budget goes to headcount, 20% to participant recruitment, and 19% to tooling — together 71% of the budget — with the remainder split across audit fees, training, and program support.[6] The majority of surveyed organizations expect GRC budgets to increase for the second consecutive year, despite a challenging economic climate.[6]**

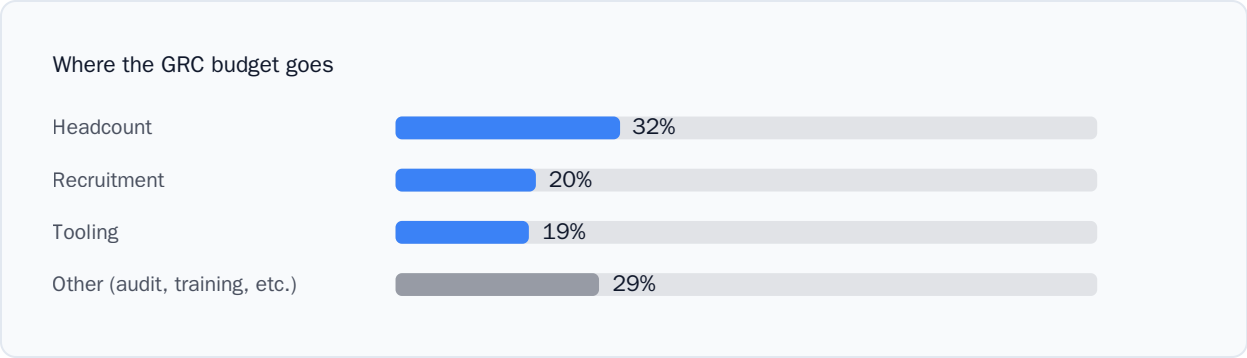

Figure 7 — Headcount, recruitment, and tooling account for 71% of GRC budgets. Source: Cycore Annual GRC Budget Survey 2025.

Headcount and labor cost

GRC labor pricing is now well-published, with figures aligning across multiple independent salary aggregators:

- **Median GRC analyst salary in the US (May 2026): approximately $95,000–$98,000 per year**, with entry-level near $70,000 and senior-level reaching $143,000+.[12] [13]
- **GRC analyst hourly rate (May 2026): $46.95**, or $97,659 annualized.[13]

These are individual contributor compensation figures. Manager, director, and CISO compensation runs materially higher and is documented in the broader cybersecurity workforce literature.

ISACA's State of Cybersecurity 2025 documents the labor environment GRC programs are hiring into: **55% of cybersecurity teams describe themselves as understaffed**, **65% have unfilled cybersecurity positions**, **70% expect demand for technical cybersecurity professionals to rise in**

the next year, and **52% of organizations report struggling to retain qualified cybersecurity staff.**[8] **66% of cybersecurity professionals report their role is more stressful now than five years ago.**[8]

Figure 8 — The labor market GRC programs are hiring into. Source: ISACA State of Cybersecurity 2025.

The implication for GRC budgeting: headcount cost is rising, fill rates are falling, and retention is harder than at any point in the past five years. The 32% of budget allocated to headcount is the floor, not the ceiling.

Tooling spend

Public benchmarks on GRC platform pricing converge on a wide range driven by company size and feature breadth:

- **Annual subscriptions for leading GRC suites: $50,000 to $500,000.**[18]
- **Implementation cost: typically 2–6× the license fee.**[18]
- **Startup-focused continuous-compliance platforms: a few thousand dollars per month** for SMB tiers, scaling into six-figure annual contracts at enterprise tiers.[18]

Gartner's published forecast suggests **legal and compliance departments will increase GRC tool investment by 50% by 2026**[18] — a meaningful inflation rate that GRC leaders should reflect in three-year budget projections.

External audit cost

SOC 2 audit pricing has been mapped repeatedly by the audit-services market, with figures that triangulate across multiple independent vendor and audit-firm sources:

- **SOC 2 Type 1: $5,000–$25,000** in CPA firm auditor fees.[11]
- **SOC 2 Type 2 from a specialist firm (A-LIGN, KirkpatrickPrice, Schellman, Prescient Security): $15,000–$75,000.**[11]

- **SOC 2 Type 2 from a regional CPA firm (Moss Adams, Withum, Aprio, Linford): $20,000–$95,000.**[11]
- **All-in mid-market first-time SOC 2 program: $60,000–$100,000+,** factoring readiness, audit, tools, penetration testing, training, and legal review.[11]
- **Internal labor for SOC 2: 100–500+ hours,** often equating to $50,000–$75,000 of project-lead time over six months.[11]

For ISO 27001, public research is less standardized on pricing but converges on the time-cost figure cited in Section 2: **roughly 150 internal hours and 80 external-assistance hours through the gap-assessment and documentation phase** for a mid-sized company.[14]

Total program economics — a synthesis

Combining headcount, tooling, and audit data, a credible mid-market multi-framework GRC program economic profile in 2026 looks roughly like:

COST DRIVER	ANNUAL RANGE (MID-MARKET, 2 FRAMEWORKS)
GRC team (1–3 FTE)	$150K – $450K
GRC platform subscription	$50K – $250K
External audit fees (2 frameworks)	$40K – $150K
Penetration testing	$20K – $60K
Training, legal, miscellaneous	$20K – $50K
Total	**$280K – $960K**

These are illustrative ranges synthesized from public salary, audit, and tooling data;[12] [13] [11] [18] they should be read as orders of magnitude rather than benchmarks. The ranges widen materially for programs running four-plus frameworks or with regulated-industry requirements (HIPAA, PCI DSS Level 1, FedRAMP) that add their own labor and audit-fee multipliers.

What we cannot yet measure (deferred to next edition)

- **Cost per framework**, isolated cleanly — i.e., the marginal cost of adding a second, third, or fourth concurrent audit on top of an established first one.
- **Cost per control** as a benchmarking metric — useful for executive reporting but rarely published.
- **Headcount-to-revenue ratios** and **headcount-to-employee-count ratios** at modal mid-market and enterprise tiers, beyond the SCCE survey aggregates.[9]

- **The cost of a failed audit** — particularly the operational drag of remediation cycles that don't show up in a year-one budget.

Cross-cutting observations

Reading across the six sections, four themes recur:

1. The center of gravity has shifted from "achieve compliance" to "operate compliance." Multi-framework is the median state. Audits per year are now 2–4 for most programs. Common controls frameworks are climbing through majority adoption. The question is no longer "how do we pass our SOC 2." It is "how do we run four overlapping audit cycles a year without burning out a small team."

2. The bottleneck is administrative work, not technical work. 76% of GRC professionals spend 30%+ of their time on repetitive admin.[7] 63% of TPRM teams are understaffed.[3] 55% of cybersecurity teams describe themselves as understaffed.[8] AI's largest GRC value is in evidence collection, control mapping, and questionnaire response — not in risk reasoning.

3. Trust gaps inside the system are widening. 81% of vendors pass questionnaires, but only 14% of TPRM teams trust the answers.[3] 60% of organizations would change auditors for quality.[1] 33% have no AI compliance strategy and 67% have neither.[1 4] The instruments the industry has used for two decades to attest control posture — questionnaires, audit reports, self-attestation — are visibly losing credibility with the people who use them.

4. AI is both the productivity unlock and the next-largest unmanaged risk. Hyperproof's 97% AI-in-workflow figure[7] and IBM's $670K shadow-AI breach premium[4] are the same phenomenon from two angles. GRC teams that succeed in 2026 will be the ones that internalize both sides.

What we are watching in 2027

Five questions this report cannot fully answer but expects to revisit:

1. **The EU AI Act's enforcement impact** — after the August 2, 2026 high-risk obligation date.[16] Specifically: how aggressively are national supervisory authorities enforcing, and what does the first wave of fines look like.

2. **The CMMC compliance curve.** A-LIGN's 94% federal-framework engagement figure[1] is a 2025–2026 leading indicator; the trailing indicator is how many programs that *started* CMMC actually *certify*.
3. **GRC tooling market consolidation.** With 100+ vendors[17] and a 50% investment increase forecast through 2026,[18] the question is which platforms emerge as durable category leaders and which segments — TPRM, AI governance, common controls — get absorbed into broader platforms.
4. **The vendor questionnaire's future.** The 14% trust figure[3] is unsustainable. Either questionnaires evolve (machine-readable, evidence-linked, continuously validated) or they get displaced by ratings, continuous monitoring, or some hybrid.
5. **GRC headcount-to-revenue benchmarks at the mid-market tier.** This is the single most-asked benchmarking question we receive and is the highest-priority data point for the next edition's primary survey.

Glossary

CMMC — Cybersecurity Maturity Model Certification. US Department of Defense framework for defense industrial base contractors.

Common controls framework (CCF) — A unified internal control set that maps to multiple external compliance frameworks (e.g., SOC 2, ISO 27001, HIPAA), enabling evidence reuse.

EU AI Act — European Union regulation on AI systems. High-risk obligations enforceable August 2, 2026.[16]

FedRAMP — Federal Risk and Authorization Management Program. US government program for cloud service authorization.

ISO/IEC 27001 — International standard for information security management systems. ISO 27001:2022 is the current version; mandatory for all new certifications.

ISO/IEC 42001 — International standard for AI management systems, increasingly cited in vendor due-diligence questionnaires.

NIS2 — EU Network and Information Security Directive 2. Drives ISO 27001 demand across EU regulated entities and supply chains.

NIST AI RMF — National Institute of Standards and Technology AI Risk Management Framework. The de facto US AI governance reference.

SOC 2 — Service Organization Control 2. AICPA framework for service-organization controls relevant to security, availability, processing integrity, confidentiality, and privacy.

Shadow AI — Use of AI tools by employees outside of formal IT or security sanction.

TPRM — Third-Party Risk Management.

About Talarity

Talarity is a GRC platform that unifies governance, risk, compliance, vendor management, and AI insights into a single operational workspace. We publish this benchmark report annually because we believe practitioners deserve a citeable, cross-cut view of where their programs stand — and because the current public research landscape is too fragmented to provide one.

For corrections, source contributions, or to participate in the primary-survey panel for the next edition, write to **research@talarity.com**.

Sources and citations

This report represents Talarity's synthesis of publicly available GRC industry research as of May 2026. Statistics are reproduced from cited primary sources; opinions and analytical framing are Talarity's own. We welcome corrections and source contributions at research@talarity.com for incorporation in the next edition.

Footnotes

1. A-LIGN, "A-LIGN Releases 2026 Compliance Benchmark Report, Unveils How Compliance Teams Can Navigate Evolving Governance Landscape," January 28, 2026. https://www.a-lign.com/articles/a-lign-releases-2026-compliance-benchmark-report — n=1,043 global respondents, fielded August–September 2025. ↩ ↩2 ↩3 ↩4 ↩5 ↩6 ↩7 ↩8 ↩9

2. Secureframe, “2026 Trends Report” findings on ISO 27001 adoption, NIS2 enforcement, and ISO 27001:2022 transition, summarized via Konfirmity, “SOC 2 Controls Mapped To ISO 27002: A 2026 Guide for Busy Teams.” https://www.konfirmity.com/blog/soc-2-controls-mapped-to-iso-27002 ↩ ↩2 ↩3 ↩4 ↩5 ↩6 ↩7 ↩8 ↩9 ↩10 ↩11
3. Atlas Systems, “120+ Third-Party Risk Management Statistics You Shouldn’t Miss,” aggregating 2025 TPRM survey research. https://www.atlassystems.com/blog/third-party-risk-management-statistics ↩ ↩2 ↩3 ↩4 ↩5 ↩6 ↩7 ↩8 ↩9 ↩10
4. IBM, “2025 Cost of a Data Breach Report” findings on shadow AI cost premium and AI governance policy adoption, summarized via Vectra AI, “Shadow AI explained: risks, costs, and enterprise governance.” https://www.vectra.ai/topics/shadow-ai ↩ ↩2 ↩3 ↩4 ↩5 ↩6
5. MarkTechPost, “Enterprise AI Governance in 2026: Why the Tools Employees Use Are Ahead of the Policies That Cover Them,” May 13, 2026, aggregating Acuvity 2025 State of AI Security, Komprise 2025 IT Survey, IBM 2025 Cost of a Data Breach, and industry skills demand data. https://www.marktechpost.com/2026/05/13/enterprise-ai-governance-in-2026-why-the-tools-employees-use-are-ahead-of-the-policies-that-cover-them/ ↩ ↩2 ↩3 ↩4 ↩5 ↩6 ↩7 ↩8 ↩9 ↩10 ↩11 ↩12
6. Cycore, “Annual GRC Budget Survey 2025 – Interactive Charts.” https://www.cycoresecure.com/blogs/annual-grc-budget-survey-2025-interactive-charts ↩ ↩2 ↩3 ↩4
7. Hyperproof, “Top 5 Takeaways from the 2026 IT Risk and Compliance Benchmark Report.” https://hyperproof.io/resource/top-5-takeaways-from-the-2026-it-risk-and-compliance-benchmark-report/ — sixth annual edition. ↩ ↩2 ↩3 ↩4 ↩5 ↩6 ↩7 ↩8 ↩9 ↩10 ↩11 ↩12 ↩13 ↩14 ↩15
8. ISACA, “State of Cybersecurity 2025 Global Press Release.” https://www.isaca.org/about-us/newsroom/press-releases/2025/state-of-cybersecurity-2025-global-press-release — n≈4,000 cybersecurity professionals worldwide. ↩ ↩2 ↩3 ↩4 ↩5
9. Society of Corporate Compliance and Ethics, “2025 Cross Industry Compliance Staff and Budget Benchmarking and Guidance Survey.” https://www.corporatecompliance.org/publications/surveys/2025-cross-industry-compliance-staff-and-budget-benchmarking-and-guidance ↩ ↩2
10. Gartner, “Global AI Regulations Fuel Billion-Dollar Market for AI Governance Platforms,” February 17, 2026 press release. https://www.gartner.com/en/newsroom/press-releases/2026-02-17-gartner-global-ai-regulations-fuel-billion-dollar-market-for-ai-governance-platforms — based on a November 2025 survey of 302 cybersecurity leaders. ↩ ↩2
11. Secureframe, “How Much Does a SOC 2 Audit Cost in 2025?” — corroborated against multiple independent vendor cost guides converging on similar ranges. https://secureframe.com/hub/soc-2/audit-cost ↩ ↩2 ↩3 ↩4 ↩5 ↩6 ↩7 ↩8
12. Salary.com, “Governance Risk And Compliance Analyst Salary,” United States, March 2026. https://www.salary.com/research/salary/position/governance-risk-and-compliance-analyst-salary ↩ ↩2 ↩3
13. ZipRecruiter, “Salary: GRC Analyst (May 2026) United States.” https://www.ziprecruiter.com/Salaries/Grc-Analyst-Salary ↩ ↩2 ↩3 ↩4
14. Konfirmity, “ISO 27001 Audit Timeline: A Practical Guide with Steps & Examples (2026).” https://www.konfirmity.com/blog/iso-27001-audit-timeline ↩ ↩2 ↩3
15. Drata, “How Much Does a SOC 2 Audit Cost?” https://drata.com/learn/soc-2/cost ↩

16. Global AI Compliance Consortium, “Global AI Governance Comparison 2026: EU AI Act vs NIST AI RMF vs ISO/IEC 42001.” https://gaicc.org/blog/ai-governance-comparison-eu-ai-act-nist-iso-42001/ ↩ ↩2 ↩3 ↩4 ↩5
17. Gartner, “Magic Quadrant for Governance, Risk and Compliance Tools, Assurance Leaders” and related “Market Guide to GRC Tools for Assurance Leaders.” https://www.gartner.com/en/documents/7111830 ↩ ↩2
18. GC Partners, “The Future of the GRC Stack: From Point Tools to Unified Risk Operating Systems,” November 4, 2025, citing GRC market structure and tooling spend forecasts. https://gcpartners.tech/2025/11/04/future-of-the-grc-stack/ ↩ ↩2 ↩3 ↩4 ↩5 ↩6 ↩7
19. Censinet, “From $2.8M Fragmented Spending to $750K Unified Platform: The Economics of Intelligent GRC.” https://www.censinet.com/perspectives/economics-intelligent-grc-unified-platform ↩