CONTINUUM GRC
Your Roadmap to Risk Reduction

August 2026
Independent customer-benchmark study
N = 275 programs

2026 AI GOVERNANCE BENCHMARK REPORT

INVENTORY, RISK & CONTROL EVIDENCE FOR AI

How 275 Continuum GRC programs govern generative AI, machine-learning systems, and third-party AI services — framework adoption, inventory completeness, evidence gaps, and the labor cost of getting AI into the compliance program.

68%
Programs with at least one AI system in scope

4.2
Average AI systems or services inventoried

19%
Pursuing ISO/IEC 42001

41%
Have complete system or model cards

EXECUTIVE SUMMARY

59%
AI systems without a documented risk classification

28%
Programs with bias or fairness testing evidence on file

180 → 55
Annual AI-governance hours, manual versus automated programs

AI is already inside most compliance scopes in this sample. What is missing is not interest — it is inventory discipline, risk classification, and assessable evidence. Programs that bolted AI onto an existing SOC 2 or NIST library without a system register, human-oversight records, and vendor AI clauses produced the largest gap lists. Programs that treated AI systems as first-class assets — with owners, risk tiers, and mapped controls — closed those gaps faster and spent less labor sustaining them.

> Labor figures use a **$125 per hour** fully loaded blended rate. They cover internal inventory, policy, risk assessment, evidence collection, and control authoring for AI systems. They exclude model development cost, external audit fees, and model-hosting spend.

STUDY METHODOLOGY

Independent analysis of 275 Continuum GRC customer programs from January 2025 through June 2026. An AI system is any generative AI service, trained or fine-tuned model, decisioning model, or third-party AI feature that processes organizational data or influences a business or security decision in scope of a formal assessment.

Of the 275 programs, 187 (68%) had at least one AI system in scope. Maturity, inventory, and evidence metrics below are calculated on that AI-active subset unless noted as full-sample. Framework mix of the full sample: CMMC 38%, FedRAMP/StateRAMP 22%, SOC 2 19%, NIST 800-53/FISMA 12%, other 9%.

AI systems in scope

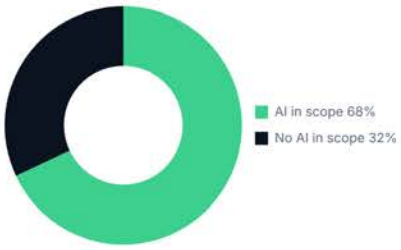

AI GOVERNANCE MATURITY

Most programs are past “no policy” and still short of continuous control. The middle band has an AI use policy and a partial inventory, but incomplete risk tiers, thin testing evidence, and weak vendor AI clauses.

Maturity bands (AI-active programs, n=187)

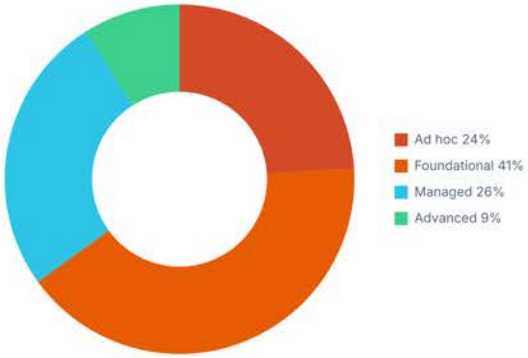

Band	Share	Typical state
Ad hoc	24%	AI in use, no register, reactive answers to assessors
Foundational	41%	Policy exists, partial inventory, limited testing evidence
Managed	26%	Full inventory, risk tiers, mapped controls, recurring review
Advanced	9%	ISO 42001 or equivalent operating, continuous monitoring, vendor AI clauses enforced

FRAMEWORK ADOPTION

ISO/IEC 42001 is rising but still early. NIST AI RMF is the most common mapping target among U.S. programs. EU AI Act obligations appear where products or users sit in the EU, not only where the company is headquartered.

AI framework posture (AI-active programs)

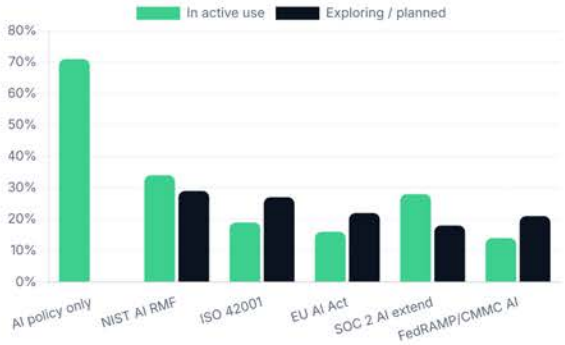

Framework or obligation	In active use	Exploring / planned
Internal AI use policy only	71%	—
NIST AI RMF mapping	34%	29%
ISO/IEC 42001 (pursue or certified)	19%	27%
EU AI Act readiness work	16%	22%
SOC 2 criteria extended for AI	28%	18%
FedRAMP / CMMC AI overlay notes	14%	21%

INVENTORY & CLASSIFICATION

You cannot govern what you have not named. Average AI-active program lists 4.2 systems or services. Completeness of that list — and a risk tier for each entry — separates managed programs from foundational ones.

Inventory completeness

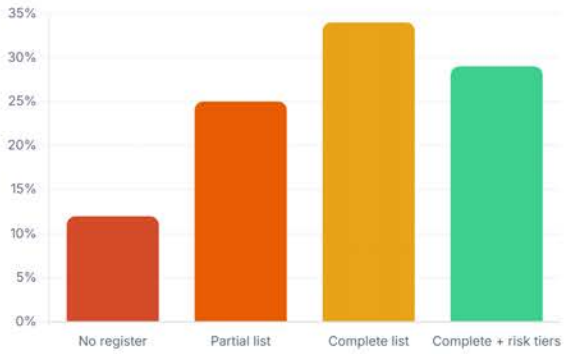

What is on the register

Asset type	Share of inventoried items
Third-party SaaS with generative AI features	38%
Enterprise copilots / assistants	24%
Internally fine-tuned or hosted models	17%
Decisioning / scoring models	12%
Other (agents, RPA+LLM, research tools)	9%

> 59% of listed systems still lacked a formal risk classification (e.g., limited / high-impact / prohibited use) at the last review snapshot.

CONTROL EVIDENCE GAPS

Assessors ask for proof of inventory, risk assessment, human oversight, data handling, testing, and incident response for AI systems. The weakest evidence families in this sample were bias/fairness testing and post-deployment monitoring records.

Share of AI-active programs with current evidence

Evidence type	Present	Coverage
Approved AI use policy	71%	
AI system inventory	63%	
Human oversight / escalation design	52%	

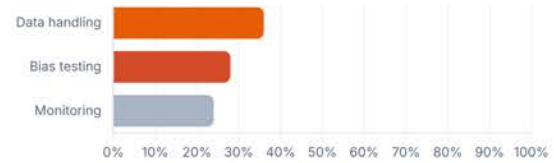

Complete system or model cards	41%
Data-handling / training-data notes	36%
Bias / fairness testing evidence	28%
Post-deployment monitoring records	24%

THIRD-PARTY AI RISK

38%
Of inventoried items are third-party SaaS AI features

46%
Third-party AI vendors with a completed AI-specific assessment

31%
Contracts with explicit AI data-use and model-training clauses

Most AI exposure in this sample arrives through vendors, not custom model training. Programs that only ran a generic SOC 2 questionnaire on those vendors missed model-training rights, retention of prompts, subprocessors, and opt-out of training. The 31% with explicit AI clauses were far less likely to carry open findings on third-party AI during assessment.

LABOR & COST

AI governance is a layer on top of existing frameworks, not a free add-on. Manual programs spent about 180 hours a year keeping inventory, risk reviews, policies, vendor AI assessments, and control narratives current. Automated programs spent about 55 hours.

180 hrs
Annual AI-governance hours, manual programs

$22,500
Manual annual labor at $125/hr

55 hrs
Annual hours when inventory, mapping, and authoring are automated

$6,875
Automated annual labor — 69% below the manual path

Annual AI-governance hours

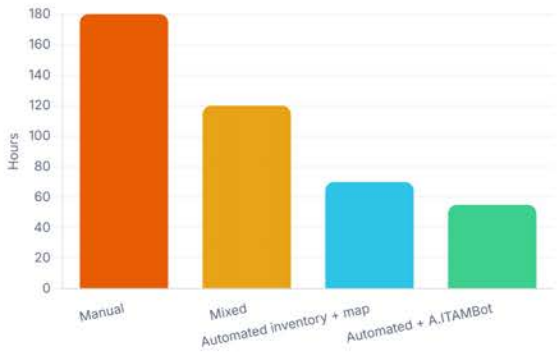

Annual AI-governance labor cost

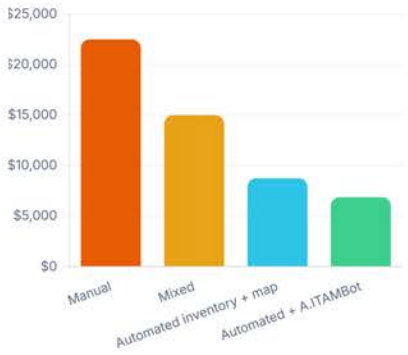

AUTOMATION & A.ITAMBOT

92%
FASTER

A.ITAMBot accelerates AI control narratives and readiness writing

Where AI controls are mapped in Continuum GRC, A.ITAMBot reduced authoring time for AI policies, system-card language, risk statements, and assessment narratives by 92% — the same rate measured on classic control writing — because the bottleneck is structured drafting from a mapped record, not the subject matter alone.

What automation covers well

- AI system register as a living

What still needs humans

- Risk-tier decisions and prohibited-use calls

Measured effect

- 69% less annual AI-governance labor when

inventory

- Control mapping from NIST AI RMF / ISO 42001 into the unified library
- Vendor AI questionnaire workflows and evidence storage
- Narrative drafts for policies, system cards, and risk statements

- Acceptance of residual model risk
- Bias testing design and interpretation
- Live demos of human oversight for assessors

inventory and mapping are automated

- 92% less time on AI-related technical writing with A.ITAMBot
- Fewer late findings on "no inventory" and "no AI policy"

RECOMMENDATIONS

1. Build the AI system register first. Everything else — risk tier, control map, vendor clause — depends on a named inventory.
2. Classify every entry (limited, high-impact, prohibited). 59% without a tier cannot prioritize testing or oversight.
3. Map AI controls into the same library as SOC 2, NIST, CMMC, or FedRAMP. Do not maintain a parallel spreadsheet program.
4. Require AI-specific vendor questions and contract language for any SaaS feature that trains on or retains prompts.
5. Collect operating evidence for human oversight and post-deployment monitoring the same way you collect access reviews — on a schedule.
6. Use A.ITAMBot for policy, system-card, and risk-statement drafts once the inventory and control map exist.

CONTINUUM GRC AND AI GOVERNANCE

Continuum GRC (IT Audit Machine®) and A.ITAMBot™ treat AI systems as governed assets: inventory, risk classification, control mapping across NIST AI RMF and ISO 42001, vendor evidence, and assessor-ready narratives. In this sample, automated AI-governance programs ran at about 55 hours a year versus 180 hours for manual programs — a 69% reduction — before counting the 92% authoring gain on AI-specific technical writing.

AI governance is no longer optional for programs that already claim strong security and privacy controls. The benchmark is whether the AI system is named, classified, evidenced, and written into the same system of record as every other in-scope control.

CONTINUUM GRC | Lazarus Alliance

Independent study. Anonymized aggregate data from 275 Continuum GRC customer programs, January 2025–June 2026.

For a program-specific AI governance comparison, contact Continuum GRC.

2026 AI Governance Benchmark Report

REQUEST A PERSONALIZED DEMO

Speak with our team using the form below or call us at 1-888-896-6207 for assistance.

Name *

First

Last

Email *

Email

Confirm Email

Phone

(201) 555-0123

Company *

How may we assist you? *

Do you want Managed Services, or Readiness Services? *

- [ ] Risk Management Services
- [ ] Compliance Management Services
- [ ] Readiness Services and Technical Writing
- [ ] Something not listed?

What type of subscription are you interested in? *

- [ ] MSP (Provider Managing Multiple Clients)
- [ ] Enterprise (Managing your Company)
- [ ] Subscribers (Limited Access to Pre-Configured Modules)
- [ ] Something not listed?

Did you want to become an integration partner?

- ○ Yes
- ○ No

What questions do you have? *

What is your hosting preference?

- ○ AWS Hosting (Commercial)
- ○ FedRAMP Authorized AWS GovCloud hosting (Federal and State Government)
- ○ AWS EU Hosting (Commercial)

Which services are you interested in? *

- [ ] GovRAMP
- [ ] FedRAMP
- [ ] SOC 1-2-3
- [ ] CMMC, NIST 800-171 & 800-172
- [ ] NIST 800-53, SCA-V
- [ ] PCI DSS QSA & SAQ
- [ ] ISO 27001
- [ ] ISO 27017
- [ ] ISO 27018
- [ ] ISO 27701
- [ ] ISO 22301
- [ ] ISO 9001
- [ ] ISO 90003
- [ ] ISO 42001
- [ ] HIPAA NIST 800-66
- [ ] CJIS
- [ ] IRS 1075
- [ ] IRS 4812

- [ ] ACAB LA DMF
- [ ] FDA 21 CFR Part 11 GxP
- [ ] MARS-E
- [ ] NIST CSF
- [ ] CBFP. FFIEC, SEC, NFA & FINRA
- [ ] C5
- [ ] Privacy, CCPA, CPRA, PIPEDA, DPIA, LGPD
- [ ] Risk Management
- [ ] COSO SOX
- [ ] NERC CIP
- [ ] CIS
- [ ] MPA Content Security Best Practices (TPN)
- [ ] FTC SafeGuards
- [ ] CTPAT
- [ ] ITAR
- [ ] EUCS
- [ ] SSDF NIST 800-218
- [ ] Something not listed

Which Governance & Policy subscription modules are you interested in? *

- [ ] GovRAMP
- [ ] FedRAMP
- [ ] SOC 1-2-3
- [ ] CMMC, NIST 800-171 & 800-172
- [ ] C5 ENS EUCS
- [ ] PCI DSS QSA & SAQ
- [ ] ISO
- [ ] HIPAA NIST 800-66
- [ ] Privacy, CCPA, CPRA, PIPEDA, DPIA, LGPD
- [ ] CJIS
- [ ] NIST 800-53
- [ ] NERC CIP
- [ ] SEC, NFA & FINRA
- [ ] Something not listed

Download our company brochure.

- [x] I agree to receiving additional marketing communications.

Success! CLOUDFLARE Privacy • Help

Submit

Looking for something?

Search Go