THOUGHT LEADERSHIP SERIES

THE ARCHITECTURE OF IMPARTIAL AUTOMATION

WHY AI REQUIRES A NEW MODEL FOR CONSULTING, INDEPENDENCE, AND CONFORMITY ASSESSMENT

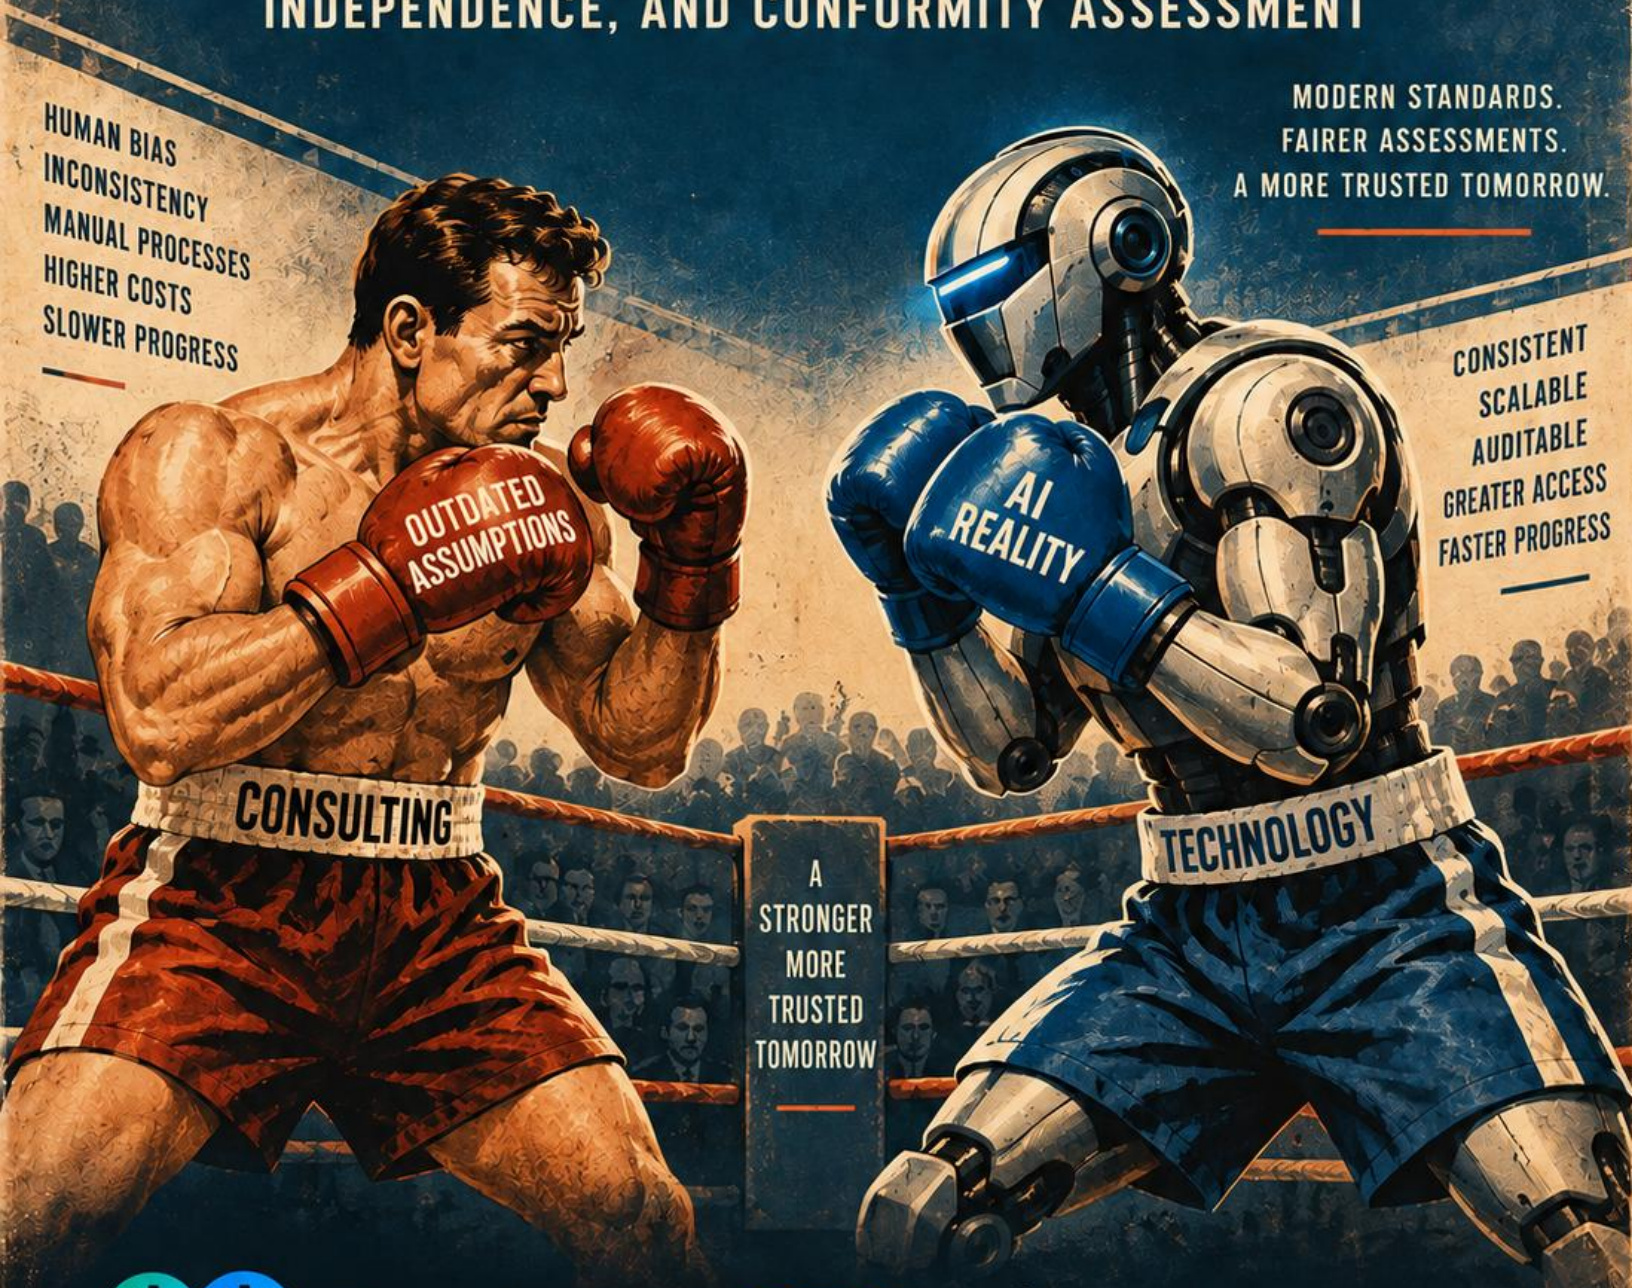

Continuum GRC®

AUTOMATE A MORE RESILIENT TOMORROW

A PUBLICATION BY
MICHAEL D. PETERS

SEPTEMBER 2026

PEOPLE | PROCESS | TECHNOLOGY | A STRONGER TOMORROW

ANALYSIS · SEPTEMBER 2026

The Architecture of Impartial Automation

Why AI Requires a New Model for Consulting, Independence, and Conformity Assessment

A policy and architecture argument for replacing person-based assumptions with functional independence: separate preparation and assessment intelligences, controlled data boundaries, auditable influence paths, and mandatory human professional judgment.

Core thesis

A software platform is not inherently a consultant. Independence rules should regulate conflicts of interest, influence, information flow, and decision authority, not automation itself. The industry must stop regulating software by pretending there is a person inside it.

Continuum GRC

FedRAMP Authorized GRC Platform · Patent-pending AI Audit Machine

Executive Summary

The conformity-assessment industry is entering a collision between rules written for human service relationships and technology that increasingly performs work once associated with human professionals. That collision is most visible when an AI-enabled platform assists an organization with readiness and also provides technically segregated capabilities that assist an independent assessor. The inherited instinct is to ask whether the software has "consulted" for the organization. This paper argues that the question is becoming conceptually wrong.

Software can perform sophisticated professional-support functions without becoming the professional. Tax software can prepare a return without becoming an accountant. Legal research software can analyze authorities without becoming counsel. A vulnerability scanner can test systems without becoming a penetration-testing firm. A GRC platform can map controls, organize evidence, draft first-pass narratives, and identify gaps without automatically becoming a consultant. Functional similarity does not establish legal or ethical identity.

That does not mean existing independence rules can be ignored. Current programs may expressly classify certain provider-owned or affiliated tools as consulting, impose disclosure requirements, or require accreditation-body or program approval. Those requirements remain controlling until changed. The purpose of this paper is not to deny them. It is to argue that the policy model behind them should evolve.

The historical rule was built around a legitimate concern: a human consultant should not design a client's compliance program and then independently certify that same work. That concern remains valid. What no longer follows is the assumption that any software performing preparation functions should inherit the same identity, incentives, memory, discretion, and conflict profile as the human consultant it may partially replace.

Continuum GRC proposes a different standard: functional independence. Under this model, the decisive questions are whether preparation can influence assessment, whether private readiness information can contaminate the assessment record, whether assessment logic can leak back into coaching, whether a human assessor retains unrestricted authority to reject automated outputs, whether organizational financial and management interests are identified and governed, and whether the entire chain of influence is auditable.

The architecture described here, separately coded preparation and assessment functions, restricted data scopes, logged boundaries, controlled promotion of evidence, human approval gates, and independent quality-system ownership, is offered as a demonstration that software can be engineered to support separation more explicitly than traditional human workflows often can. The architecture does not erase existing program rules. It shows why those rules should be rewritten around actual influence rather than analogies to people.

The policy challenge

Rules designed to stop a consultant from auditing the consultant's own advice remain necessary. Extending those rules mechanically to technically segregated software is not. The next generation of independence policy should distinguish human service conflicts from controlled automated capability.

I. The Problem AI Was Built to Solve Is Not Speed. It Is Human Variability.

Speed is the easiest benefit to sell and the least important one to understand. The deeper opportunity is consistency. Human review is indispensable, but human beings vary: they tire, anchor on early impressions, remember recent failures, interpret ambiguous evidence differently, and carry different levels of framework knowledge into an engagement. Mature quality systems compensate through training, calibration, workpapers, supervision, peer review, conflict controls, and independent quality assurance. AI can add another layer: repeatable processing that does not fluctuate simply because the assessor is on control 87 instead of control 7.

That is not an argument that AI is unbiased. AI can reproduce bad assumptions, encode poor rules, rely on incomplete retrieval, misclassify evidence, or create automation bias when users over-trust an apparently precise output. The relevant comparison is not "biased human versus unbiased machine." It is whether a governed combination of automation and professional judgment can reduce known forms of variance while making new machine risks visible and controllable.

Human bias is a quality risk to manage, not a moral judgment.

Confirmation bias can favor evidence that supports an emerging narrative. Anchoring can give disproportionate weight to early impressions. Availability bias can overemphasize a recent incident or memorable interview. Familiarity can soften skepticism. These are predictable characteristics of human cognition. An assessment system that pretends they do not exist is weaker than one that designs controls around them.

A configured automation engine can apply the same evidence schema, crosswalk rules, freshness checks, completeness criteria, and provenance tests repeatedly. That can reduce specific forms of human inconsistency. It should not make final conformity decisions. Instead, it should produce traceable work for qualified professionals to accept, reject, modify, and explain.

Fatigue and clerical omission are engineering problems.

Long assessments involve repetitive artifact review, interview notes, sampling, timestamps, version checks, inheritance records, remediation status, and cross-reference work. Human attention is not constant across days or weeks. The industry traditionally addresses this by adding labor and review. Automation offers a second approach: remove repetitive processing from the critical path and preserve scarce professional attention for matters that actually require skepticism and judgment.

No organization should claim that AI eliminates error. The stronger claim is narrower: a well-designed system can make selected errors easier to prevent, detect, reproduce, and audit.

II. Two Jobs, Two Intelligences: Preparation Is Not Assessment

A compliance lifecycle contains two fundamentally different jobs. The organization prepares. The assessor evaluates. The existence of software on both sides does not collapse those jobs into one role any more than both parties using email makes them one organization.

The preparation function

The organization must determine scope, implement controls, document the environment, collect evidence, track deficiencies, remediate gaps, and assemble a package. A preparation intelligence can help interpret requirements, map overlapping frameworks, identify missing artifacts, draft first-pass narratives from

structured information, detect stale evidence, and maintain readiness over time. Management still owns the facts, the implementation, and what is ultimately submitted.

The assessment function

The assessor must test the submitted package against applicable criteria, perform or direct sampling, conduct interviews, evaluate contradictory evidence, determine findings, apply professional skepticism, and issue the conclusions assigned to qualified personnel or the responsible assessment organization. A separately coded assessment intelligence can assist with requirement tracing, completeness checks, consistency analysis, workpaper structure, and analytical flags. It should not remediate the client, coach the inspectee during assessment, conceal adverse evidence, or issue the final professional conclusion.

The distinction that matters is not whether both functions are software. The distinction is whether one can improperly influence the other.

Side-by-side: what should remain separated

Dimension	Preparation intelligence	Assessment intelligence
Purpose	Help the organization build a complete, current, reviewable package	Help qualified assessors process the submitted package against applicable criteria
Primary user	System owners, GRC teams, program managers	Qualified assessors, reviewers, inspection personnel
Typical work	Mapping, evidence organization, draft narratives, gap and remediation tracking	Requirement tracing, completeness checks, inconsistency flags, workpaper support
Data scope	Organization workspace, connected systems, drafts controlled by the owner	Officially submitted evidence and assessor workspaces
May not do	Issue assessment conclusions or conceal adverse facts	Remediate the client, coach mid-assessment, or issue conclusions without required human authority
Human gate	Organization owner reviews and accepts generated content	Qualified assessor reviews, adjusts, rejects, or approves material outputs

III. The Category Error: A Software Platform Is Not Inherently a Consultant

The most consequential assumption this paper challenges is the equation of software assistance with consulting. That equation emerged from an era in which capability and professional identity were tightly coupled: if someone analyzed your controls, drafted your documentation, and recommended remediation, that "someone" was a person or firm. AI breaks that coupling.

Software may now perform tasks that resemble professional services without acquiring the legal identity, judgment, incentives, loyalties, memory, or discretion of the professional who historically performed them. The industry therefore needs a more precise test than "the tool helped the customer."

A more precise proposition

A software platform is not inherently a consultant. Software may perform functions historically performed by consultants, but functional similarity does not establish identity. The conflict analysis should focus on who controls decisions, who can influence outcomes, what information crosses boundaries, what economic interests exist, and whether professional authority remains independent.

Assistance is not the same as discretionary influence.

A system can identify that evidence is missing without deciding how management should implement the control. It can generate a first-pass SSP narrative without deciding whether that narrative is true. It can compare evidence to a requirement without deciding whether the final evidence is sufficient. It can recommend that an artifact appears stale without choosing management's remediation strategy. These are capabilities. Whether they amount to consulting should depend on the facts of human control and influence, not merely on the fact that automation made the work easier.

The analogy is already familiar elsewhere.

Tax software can calculate a tax position without becoming a CPA firm. Legal research platforms can summarize cases without becoming the attorney of record. CAD systems can generate designs without becoming licensed engineers. Vulnerability scanners can identify weaknesses without becoming the independent penetration-testing firm. In each case, sophisticated software supports a professional or management function while responsibility remains with a person or legal entity. GRC automation should be analyzed with the same conceptual discipline.

The industry must stop regulating software by pretending there is a person inside it.

Person-based conflict rules are essential when people or firms have divided loyalties, commercial pressure, prior advisory involvement, compensation incentives, or the ability to shade an outcome. But software architecture permits controls that human organizations cannot implement with the same precision: hard data boundaries, role-restricted memory, deterministic promotion of evidence, immutable logs, separate release processes, different toolchains, and machine-verifiable access policies. Policy should recognize those differences instead of erasing them.

IV. Current Rules Should Be Followed, and Then Changed

A disruptive policy argument is strongest when it states the current rule accurately. Some existing programs treat provider-owned or affiliated preparation technology as consulting or subject it to heightened impartiality review. Those requirements govern today's engagements until the relevant program, accreditation body, or rulemaker changes them. Continuum GRC does not argue that regulated parties should disregard current requirements.

The argument is that the rulemaking premise deserves reconsideration. A policy developed to prevent an inspection body from selling advisory services and then certifying its own advice can become overbroad when it treats independently operated software as though the software developer's personnel personally performed the client's preparation work.

FedRAMP illustrates the tension.

Published FedRAMP 3PAO guidance addresses tools owned or developed by a 3PAO and tools owned or developed by an organization affiliated with the 3PAO or its management. Depending on the relationship and use case, the guidance can treat direct support to a CSP as consulting and can require documented impartiality rationale and approval. That is the current framework and must be respected where applicable. See Reference 1.

But the policy question for the next generation is different: should an enterprise software product used independently by a customer be treated the same as personnel from the assessment organization advising that customer? If the software is commercially available, customer-controlled, technically segregated from assessment functions, incapable of dictating the assessment result, and subject to auditable boundaries, a blanket equivalence may preserve an old labor model rather than protect impartiality.

CMMC presents the same conceptual challenge.

CMMC correctly restricts consulting-then-assessing relationships and protects the autonomy of qualified assessment personnel. Those protections should remain. The modernization opportunity is to distinguish a human or legal entity that actually consulted for the OSC from a software capability the OSC independently used. The decisive issue should be whether prohibited advisory influence occurred and whether assessment independence can be demonstrated, not whether the same product family contains tools for both sides of the lifecycle. See Reference 3.

AICPA and ISO independence concepts likewise focus on threats, safeguards, responsibilities, relationships, and professional judgment. The policy evolution proposed here is therefore not abandonment of independence. It is a more exact way to identify what independence is supposed to protect.

V. A Replacement Doctrine: The Functional Independence Model

If the industry is going to stop equating software with consulting, it needs a replacement test. Continuum GRC proposes the Functional Independence Model: use of common technology should not, by itself, establish a consulting relationship or independence impairment. Instead, the assessment should examine the actual channels through which preparation could influence assessment.

Five tests for functional independence

Test	Question	Evidence of control
1. Decision independence	Who can determine or alter the assessment conclusion?	Qualified personnel retain authority; automated outputs are reviewable and rejectable; no auto-certification.

Test	Question	Evidence of control
2. Data independence	Can private readiness activity contaminate the assessment record?	Assessment sees only officially submitted evidence and authorized data; readiness drafts and coaching traces remain segregated.
3. Process independence	Can preparation logic manipulate assessment logic or vice versa?	Separate code paths, prompts/rules, permissions, release controls, and restricted information flows.
4. Personnel and economic independence	Do people, ownership, compensation, or commercial relationships create prohibited influence?	Conflict review, disclosures, contractual controls, governance, recusal, and program approvals where required.
5. Auditability	Can a reviewer reconstruct what the machine and humans did?	Logs, provenance, version history, overrides, approvals, model/rule versions, and quality-review records.

This model does not create a loophole.

If personnel from an assessment body actually provide prohibited consulting, functional independence fails. If preparation data secretly trains or conditions the assessment engine, it fails. If financial incentives can alter scoring, it fails. If the assessor is required to accept machine conclusions, it fails. If a program expressly prohibits the relationship, the engagement remains prohibited until the rule changes. The model is demanding precisely because it replaces a simple label with testable evidence.

What it rejects is only the automatic inference that capability equals consulting. A tool should be judged by what it does, who controls it, what information it can access, and how it can influence the outcome.

VI. The Continuum GRC Control Architecture: Separation, Boundary, Human Command

Continuum GRC's architecture is intended to demonstrate that functional independence can be engineered rather than merely promised. The relevant controls are not branding distinctions. They are technical and operational boundaries that can be inspected, tested, logged, and challenged.

Separately coded intelligences

Preparation and assessment capabilities are implemented as distinct coded functions rather than a single undifferentiated conversational agent with two labels. Separate implementation allows different instructions, permissions, tools, data scopes, evaluation methods, release cycles, and monitoring. A change designed to make readiness assistance more useful should not silently rewrite the logic used to support assessment processing.

The controlled boundary

Preparation content should cross into the assessment environment only through an explicit submission or authorized evidence path. Drafts, private working notes, exploratory prompts, rejected versions, and internal

readiness discussions should not become invisible context for the assessor's automation. Assessment methods, sampling logic, and protected evaluation criteria should likewise be prevented from leaking into the preparation layer in ways that allow a customer to game the assessment.

Human command

Every material machine output that could affect an assessment conclusion should remain subordinate to qualified human review. The assessor must be able to reject a mapping, change a risk interpretation, demand additional evidence, expand testing, document disagreement with automation, and reach a conclusion the machine did not predict. Human-in-the-loop is meaningful only when the human retains actual authority rather than serving as a ceremonial approver.

Evidence integrity and provenance

Separation means little if evidence can be silently substituted or its provenance cannot be reconstructed. Cryptographic hashing, version records, source metadata, controlled ingestion, timestamps, and audit logs support confidence that the evidence reviewed is the evidence submitted. Integrity is not independence, but independence is difficult to defend when the object under inspection cannot be reliably identified.

Governed AI change

The AI layer itself must be governed. Model versions, prompts, rules, retrieval sources, scoring logic, tool permissions, and release changes should be subject to change control proportionate to their effect on assessment work. An impartiality architecture that can be changed informally by product updates is not an architecture; it is an aspiration.

VII. What This Changes for Organizations Preparing for Assessment

The practical benefit of preparation AI is not that management stops thinking. It is that scarce experts stop spending so much time on clerical coordination and can spend more time on implementation quality.

- Control mapping can become a maintained system relationship instead of a spreadsheet exercise repeated for every framework.
- Evidence collection can move from a seasonal document hunt toward continuous, provenance-aware ingestion from connected systems.
- First-pass narratives can be generated from structured facts, while management remains responsible for making them accurate and organization-specific.
- Gaps can surface earlier, allowing remediation before the formal assessment window.
- Reused evidence can be evaluated against multiple requirements without assuming that one artifact automatically satisfies every framework.
- Preparation records can remain under the organization's control until they are intentionally submitted for independent assessment.

The policy implication is important. When an organization independently operates software to perform this work, the industry should not automatically treat the software developer as having personally designed the organization's controls. That inference should require evidence of actual advisory control or prohibited influence.

VIII. What This Changes for Independent Assessors

Assessment organizations are constrained by qualified time. They should not waste credentialed attention on every task simply because that task historically required human labor. Completeness checking, crosswalk validation, stale-evidence detection, structured workpaper assembly, and comparison of narrative against configuration are candidates for controlled automation when the method is validated and human responsibility remains clear.

The assessor's role becomes more, not less, important when automation is introduced. Someone must challenge the machine, investigate anomalies, test scope, evaluate exceptions, conduct interviews, determine whether evidence is persuasive, and own the professional conclusion. Good automation should move qualified personnel toward those responsibilities instead of trying to eliminate them.

The industry should therefore distinguish between automation that supports an assessor and automation that captures the assessor. A system crosses the line when it makes the human subservient, conceals how conclusions were produced, prevents meaningful override, or imports undisclosed preparation context. A system supports independence when it exposes its inputs, preserves override authority, logs decisions, and keeps preparation influence outside the assessment boundary unless deliberately submitted.

IX. A Policy Agenda for AI-Native Conformity Assessment

If regulators, accreditation bodies, and assessment schemes want both independence and modernization, they should replace blanket software-as-consulting presumptions with risk-based requirements that are specific enough to audit.

- Define consulting by discretionary advisory influence and management participation, not merely by the presence of automated assistance.
- Distinguish commercially available software operated by the organization from human professional services performed by the assessment body or its personnel.
- Require disclosure of ownership, affiliation, compensation, and data-flow relationships that can create real influence over an assessment.
- Permit common technology where preparation and assessment functions satisfy defined separation, access-control, logging, and human-authority criteria.
- Require documented validation of AI functions used in assessment work, including known limitations, change control, and override mechanisms.
- Prohibit hidden cross-context learning or data leakage from readiness activity into assessment decision support unless the material was deliberately submitted as evidence.
- Preserve existing restrictions on human consulting-then-assessing relationships unless and until a program affirmatively changes them.
- Create accreditation guidance for evaluating AI-enabled assessment architectures so approval is based on evidence rather than analogy.

These changes would not weaken impartiality. They would make impartiality more testable. A rule that says "do not use the same software" is simple, but simplicity is not the same as rigor. A rule that says "demonstrate independent decision authority, segregated data, controlled influence, organizational conflict governance, and auditable provenance" is harder to satisfy, and much closer to the actual risk.

X. Continuum GRC Is Not Previewing the Future. It Is Arguing for the Rules the Future Requires.

The GRC market is moving from systems of record to systems that actively interpret requirements, collect evidence, generate work product, identify anomalies, and assist professional review. That transition is already occurring. The policy choice is whether conformity-assessment rules will govern these systems according to their actual architecture or continue to map every automated capability back onto a human-services analogy.

Continuum GRC's position is deliberately disruptive: independence rules should regulate conflicts of interest, not automation itself. Where software can demonstrate technical segregation, restricted information flow, independent human authority, organizational conflict controls, and auditable decision provenance, regulators should have a path to recognize that architecture as compatible with impartial assessment.

That position does not ask the industry to trust AI. It asks the industry to demand evidence from AI systems in the same way it demands evidence from people and organizations. The difference is that software can often make its control boundaries more explicit: permissions can be tested, logs can be inspected, data paths can be traced, model versions can be recorded, and overrides can be reconstructed.

The proposition for industry debate

Rules developed to prevent a consultant from auditing the consultant's own advice remain necessary. Extending those rules mechanically to technically segregated software is not. Independence should be assessed according to influence, control, information flow, professional authority, economic interests, and demonstrable separation.

XI. Closing Position: Stop Regulating Software by Analogy

The conformity-assessment industry is right to protect independence. It is wrong to assume that protecting independence requires preserving every historical distinction created by human labor models.

A person who advises a client and then audits that advice may have a conflict. A legal entity that profits from both sides of an engagement may have a conflict. Shared management or compensation arrangements may create a conflict. Those threats are real and should remain subject to strict controls.

But a software capability is not a person. A platform does not acquire divided loyalty merely because it can perform tasks that consultants once performed manually. The relevant question is whether the technology, the people around it, or the organizational relationships can improperly influence the independent assessment.

The industry should therefore move from identity-based presumptions to evidence-based functional independence. Ask who controls the decision. Ask what information crosses the boundary. Ask whether the assessor can reject the machine. Ask whether private preparation activity can contaminate assessment. Ask who owns the technology and how financial interests are governed. Ask whether every material influence can be reconstructed after the fact.

If those controls fail, the technology should not be used. If a current program prohibits the relationship, the rule must be followed until it is changed. But where the controls can be demonstrated, the existence of a common software platform should not automatically be treated as proof of consulting or impaired independence.

The next generation of conformity assessment will not be fully autonomous, and it should not be. It will be AI-assisted, continuously evidenced, machine-verifiable in places, and governed by humans who retain professional authority. Independence frameworks need to evolve for that reality.

Two intelligences. A controlled boundary. Independent human judgment. Auditable separation.

That is not an erosion of impartiality. It is an opportunity to engineer impartiality more deliberately than the industry has ever been able to do with people alone.

References and Primary Authorities

1. FedRAMP, 3PAO Readiness Assessment Report Guide, including current guidance addressing 3PAO independence, consulting, and tools owned or developed by a 3PAO or affiliated organization. https://www.fedramp.gov/resources/documents/3PAO_Readiness_Assessment_Report_Guide.pdf

2. FedRAMP, current program, Marketplace, and FedRAMP 20x materials. https://www.fedramp.gov/

3. 32 CFR Part 170, CMMC Program requirements, including conflicts of interest, Code of Professional Conduct, ethics, and assessment responsibilities. https://www.ecfr.gov/current/title-32/subtitle-A/chapter-I/subchapter-D/part-170

4. AICPA & CIMA, Auditor Independence Resource Center and AICPA Code of Professional Conduct independence resources. https://www.aicpa-cima.com/resources/toolkit/auditor-independence-resource-center

5. ISO/IEC 17020:2012, Conformity assessment, Requirements for the operation of various types of bodies performing inspection. https://www.iso.org/standard/52994.html

6. NIST AI Risk Management Framework (AI RMF), governance and risk-management concepts relevant to trustworthy use of AI systems. https://www.nist.gov/itl/ai-risk-management-framework

About the Author

Michael D. Peters is CEO and Founder of Continuum GRC and the architect of the IT Audit Machine® and A.ITAM™ platforms. He has more than two decades of experience in cybersecurity, governance, risk, compliance, conformity assessment, and audit automation. His work has focused on developing operating models and technologies that improve evidence quality, assessment consistency, continuous compliance, and professional decision support while preserving human responsibility for material judgments and conclusions.

About Continuum GRC

Continuum GRC is an enterprise SaaS platform for governance, risk, and compliance programs. Built on the IT Audit Machine® (ITAM) and A.ITAM™ frameworks, the platform provides capabilities for evidence collection, control mapping, risk scoring, dashboards, continuous monitoring, and AI-assisted preparation and assessment workflows through A.ITAMBot™. The platform is designed around the premise that automation should increase traceability and consistency while preserving the decision rights and professional responsibilities assigned to people and organizations by applicable assessment schemes.

A.ITAMBot is the intelligent-assistance layer within A.ITAM. It is designed to automate or assist with selected routine GRC activities across supported frameworks. Actual time savings, workload reduction, accuracy, and other performance outcomes vary by environment and use case and should be represented publicly only when supported by documented measurement. Generated work product requires appropriate human review, and assessment or certification decisions remain the responsibility of qualified personnel and organizations operating under the applicable program and quality system.

IMPORTANT NOTICE

This publication is a technology-policy argument intended to encourage modernization of industry rules. It describes Continuum GRC's architecture and the author's analysis of automation in governance, risk, compliance, and conformity-assessment workflows. It does not state that current program requirements may be disregarded, and it is not legal advice, an accreditation determination, a certification decision, a guarantee of compliance, or a representation that use of any technology satisfies the independence or impartiality requirements of a particular scheme. Current program, accreditation, and professional requirements remain controlling unless and until they are changed. Organizations and assessment bodies remain responsible for evaluating conflicts of interest, impartiality threats, organizational relationships, tool usage, applicable professional standards, and program-specific requirements, including disclosure or approval obligations where applicable.