The RIA's AI Governance Playbook

What SEC Examiners Are Asking in 2026

Dr. Leigh Coney

Founder & Principal Consultant, WorkWise Solutions

Q2 2026

workwisesolutions.org

leigh@workwisesolutions.org

Abstract

The SEC has no artificial intelligence rule for investment advisers and will not adopt one soon. The single proposal it had, the 2023 predictive data analytics rule, was withdrawn in June 2025, and the pipeline behind it is empty. Advisers who read that as a reprieve have drawn the wrong conclusion. AI governance became an examination subject without ever becoming a rule. The Division of Examinations' fiscal year 2026 priorities commit staff to reviewing the accuracy of AI representations and the adequacy of the policies and procedures that supervise AI use. Examination request lists have included AI inventories, acceptable use policies, vendor diligence files, and evidence of human review since late 2023. And the amended Regulation S-P reached its final compliance date on June 3, 2026, which makes incident response and oversight of AI vendors enforceable obligations for every SEC-registered adviser, down to the smallest firm.

This paper sets out what examiners ask for, organized as seven artifacts a firm should be able to produce on request; the existing rules that make those artifacts mandatory in the absence of any AI rule; the shadow AI problem that undermines all seven; and a 90-day sequence for building the program. It closes with the behavioral design choices that determine whether oversight operates in practice or only on paper.

Contents

1. Introduction: The Exam Arrived Before the Rulebook

On June 3, 2026, the compliance clock on the amended Regulation S-P ran out for the last group of covered firms: SEC-registered investment advisers with under $1.5 billion in regulatory assets under management. Every registered adviser in the country now owes regulators a written incident response program, customer breach notification within 30 days, and documented oversight of the service providers that touch customer information, a category that includes nearly every AI tool an advisory firm uses (SEC, 2024b).

That deadline is the hardest date AI governance has ever had, and it arrived without an AI rule attached to it. The Commission's one attempt at AI-specific rulemaking, the predictive data analytics proposal of July 2023, drew heavy industry opposition and was withdrawn in June 2025 along with thirteen other pending proposals (SEC, 2023; SEC, 2025a). Nothing has replaced it. John O'Connell of The Oasis Group has made the structural case that nothing will for years: a short-staffed Commission, a chairman deregulatory by conviction, and Administrative Procedure Act timelines that put any new adviser AI rule near the end of the decade (O'Connell, 2026).

That analysis is correct about rulemaking and dangerously incomplete as a risk assessment. The examination program did not wait for a rule. The Division of Examinations began sweeping advisers' AI use in December 2023, requesting tool inventories, governance policies, testing results, and incident reports (ACA Group, 2023). The Commission brought its first AI enforcement actions against advisers in March 2024 (SEC, 2024a). The fiscal year 2026 examination priorities, published in November 2025, name AI representations and AI supervision as subjects for the current cycle (SEC, 2025c). Compliance consultants report the same requests appearing in routine examinations of ordinary wealth management firms, not just sweeps of quantitative managers (O'Connell, 2026).

An examination is not an intelligence test. It is a records test. Examiners do not grade the sophistication of a firm's AI; they grade whether the firm can show what it uses, who approved it, what data it touches, and who checked the output before a client saw it. The firms that fail are rarely doing anything exotic. They simply cannot produce the records.

The practical question for an RIA principal has changed from *what will the rule require* to *what will the examiner ask for, and can we hand it over*. This paper answers the second question. The sections that follow trace how examination replaced rulemaking, measure the distance between AI adoption and AI governance in the current survey data, specify the seven artifacts examiners request and the rules that anchor each one, address the shadow AI problem, and lay out a 90-day build. The closing sections turn to the behavioral design choices that keep oversight real after the binder is assembled.

2. Two Years of Signal: How Examination Replaced Rulemaking

The first signal was an enforcement action, not a rule. On March 18, 2024, the Commission settled charges against two advisers, Delphia (USA) Inc. and Global Predictions Inc., for false and misleading statements about their use of AI. Delphia had claimed for years, in SEC filings, a press release, and on its website, that its AI and machine learning incorporated client data into its investment process; no such capability existed. Global Predictions called itself the "first regulated AI financial advisor" and advertised AI-driven forecasts it did not produce. The firms paid $225,000 and $175,000 respectively, $400,000 in combined civil penalties (SEC, 2024a). The legal basis was entirely conventional: the Marketing Rule and the antifraud provisions of the Advisers Act. The Commission did not need an AI rule to bring an AI case, and said so plainly at the time.

Two months later came the structural change. On May 16, 2024, the Commission adopted the first substantive amendments to Regulation S-P since the rule was written in 2000 (SEC, 2024b). Covered institutions, including all registered advisers, must now maintain a written incident response program designed to detect, respond to, and recover from unauthorized access to customer information; notify affected individuals within 30 days; oversee service providers with access to customer information; and keep records evidencing all of it. Compliance split by size: December 3, 2025 for advisers at or above $1.5 billion in assets under management, June 3, 2026 for everyone else.

FINRA stated the underlying principle in plain text that June. Regulatory Notice 24-09 reminded member firms that its rules are technologically neutral: supervision, communications, and recordkeeping obligations apply to generative AI exactly as they apply to any other tool, and no new rule is required to make them bite (FINRA, 2024). The notice binds broker-dealers and dual registrants directly, but it described the posture both regulators now share.

Then the rulemaking track closed. On June 12, 2025, the Commission formally withdrew fourteen pending proposals, the predictive data analytics rule among them (SEC, 2025a). Parts of the industry read the withdrawal as retreat. The agency's next move pointed the other way: on August 1, 2025, it created an AI task force and appointed its first Chief AI Officer, Valerie Szczepanik, to put AI to work inside the agency's own operations (SEC, 2025b). A regulator that automates its own document review reads more of what firms produce, not less.

The examination priorities for fiscal year 2026, published November 17, 2025, made the program explicit. Chairman Paul Atkins framed the release as transparency, noting that examinations "should not be a 'gotcha' exercise" (SEC, 2025c). The priorities themselves commit staff to reviewing registrants' AI representations for accuracy and assessing "whether firms have implemented adequate policies and procedures to monitor and/or supervise their use of AI technologies," across functions from fraud prevention to back-office operations to trading. The cybersecurity section adds a focus on the training and security controls firms use against AI-associated risks, and the staff committed to engaging firms on Regulation S-P incident response readiness during examinations ahead of the compliance dates (SEC, 2025c).

Two December publications completed the picture. A Division risk alert on the Marketing Rule documented a recurring pattern: advisers whose written policies reflected the rule and whose practices did not, testimonials missing required disclosures, third-party ratings used without due diligence (SEC, 2025d). And FINRA's 2026 Annual Regulatory Oversight Report carried its first standalone generative AI chapter, extending the supervision discussion to autonomous AI agents and to human-in-the-loop review of model outputs (FINRA, 2026).

Date	Action	Why it matters
Dec 2023	Examinations staff open a sweep of advisers' AI use	Request lists include AI inventories, governance policies, testing results, and incident reports (ACA Group, 2023)
Mar 18, 2024	First AI-washing enforcement: Delphia and Global Predictions	$400,000 in combined penalties under existing antifraud and marketing provisions (SEC, 2024a)
May 16, 2024	Regulation S-P amendments adopted	Incident response programs, 30-day breach notification, service provider oversight (SEC, 2024b)
Jun 27, 2024	FINRA Regulatory Notice 24-09	Existing rules apply to generative AI; no new obligations needed (FINRA, 2024)
Jun 12, 2025	Predictive data analytics proposal withdrawn	The rulemaking path closes; the existing-authority path is confirmed (SEC, 2025a)
Aug 1, 2025	SEC AI task force formed; first Chief AI Officer named	The agency builds internal AI capacity of its own (SEC, 2025b)
Nov 17, 2025	Fiscal year 2026 examination priorities published	AI representations and AI supervision named as examination subjects (SEC, 2025c)
Dec 3, 2025	Regulation S-P compliance date, larger advisers	Firms at or above $1.5 billion in assets under management (SEC, 2024b)
Dec 16, 2025	Marketing Rule risk alert	Staff document divergence between written policy and actual practice (SEC, 2025d)
Jun 3, 2026	Regulation S-P compliance date, smaller advisers	The full adviser population is now covered (SEC, 2024b)

Table 1. The regulatory record on adviser AI, December 2023 to June 2026.

Read in sequence, the record is not ambiguous. The Commission chose enforcement, examination, and existing authority over new rulemaking, and it has now run that strategy through a full cycle: sweep, enforcement, amended rule, priorities, risk alert, deadline. The open question for an adviser is no longer whether the rules reach AI. It is whether the firm can evidence compliance with rules that already do.

3. The Numbers: Adoption Outran Governance

Schwab Advisor Services surveyed 533 RIAs in October 2025 and published the results in January 2026. Sixty-three percent of firms use AI tools in some capacity, more than double the share in 2023. Among users, 82 percent rely on generative AI tools, most often through individual experimentation rather than firm-wide systems. Only about one user in ten has integrated AI into business strategy (Charles Schwab, 2026). Expectations run well ahead of that: 59 percent of advisors expect AI to have a direct, measurable effect on client relationships within a year, and 68 percent expect it to reshape financial advice within three.

Compliance officers see the same picture from the other side of the desk. The 20th annual Investment Management Compliance Testing Survey, fielded to 577 adviser firms in 2025, found 57 percent naming AI the year's hottest compliance topic, the first time AI has topped the list. Off-channel communications, the previous leader, collapsed from 59 percent to 23 percent (IAA, ACA Group, & Yuter Compliance Consulting, 2025).

The same survey measured what firms have built underneath the worry, and the answer is: less. Sixty-four percent of firms have policies and procedures governing employee AI use, which leaves roughly a third without them after two years of regulatory signal. Thirty-eight percent have an AI committee or governance group. Among firms that use AI tools, 44 percent have never tested or validated the output. Asked why they have not formally adopted AI, firms cite cybersecurity and privacy concerns (60 percent) and regulatory uncertainty (52 percent) ahead of every other reason (IAA et al., 2025). The industry is anxious about precisely the things it is not yet documenting.

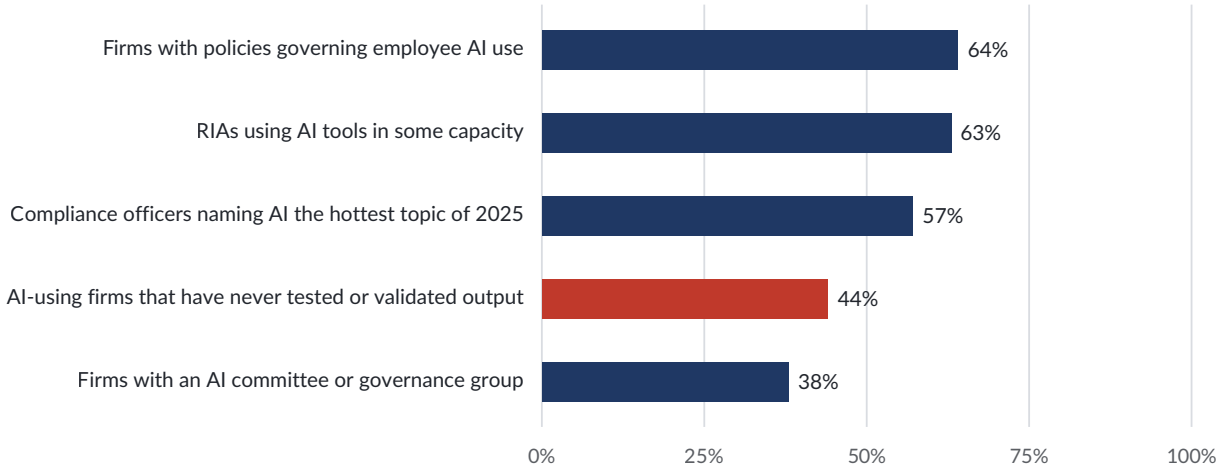

Figure 1. Adoption has outrun governance. Share of RIAs using AI tools in some capacity; share of adviser compliance professionals naming AI the hottest compliance topic of 2025; share of adviser firms with policies governing employee AI use; share of AI-using firms that have never tested or validated tool output; share of firms with an AI committee or governance group.
Sources: Charles Schwab (2026); IAA, ACA Group, & Yuter Compliance Consulting (2025).

Set those columns side by side and the examination exposure writes itself. A majority-adopted technology is a minority-governed one. Tools reach client data through individual experimentation, which means outside whatever policy exists. And the single most common control failure, untested output, sits exactly where the fiscal year 2026 priorities point: whether advice and recommendations produced with automated tools are consistent with disclosures, client profiles, and regulatory obligations (SEC, 2025c).

The deficit is not a deficit of intent. Firms are plainly worried. It is a deficit of evidence, and examinations trade exclusively in evidence.

4. The Seven Artifacts: What Examiners Ask For

The fiscal year 2026 priorities describe four assessments examiners make of firms using automated tools and AI: whether representations are fair and accurate; whether operations and controls match the disclosures made to investors; whether algorithms produce advice consistent with clients' investment profiles and stated strategies; and whether controls confirm that automated recommendations meet regulatory obligations, with particular attention to retail and older investors (SEC, 2025c).

Behind that language sits a document request. The sweep letters served on advisers since December 2023 asked for tool inventories, AI policies and the written guidance given to employees, security controls protecting client data inside AI systems, data sources and their providers, validation and testing reports, internal reports of AI incidents, AI-referencing marketing materials, governance committee lists, and continuity plans for AI system failures (ACA Group, 2023). Routine examinations of wealth management firms now carry the core of the same list (O'Connell, 2026). Across the request lists, seven artifacts recur. Together they are the playbook.

1. The AI inventory. A current record of every AI tool that touches client data, recommendations, or client communications: standalone tools, AI embedded inside vendor platforms (CRM assistants, meeting notetakers, email and proposal generators), tools used by affiliates and service providers, and the unapproved tools staff run on personal devices. The inventory comes first because every other artifact is defined against it. A policy cannot govern tools the firm has not found, and a supervision procedure cannot review outputs the firm does not know exist.

2. The acceptable use policy. A written document that classifies tools as approved, restricted, or prohibited; specifies which categories of client data are permitted in which class of system; names the person who approves new tools and the turnaround the firm commits to; and carries an update cadence matched to the pace of the tools, annual at an absolute minimum and quarterly in practice (O'Connell, 2026). A policy dated before the firm's current tool population reads as evidence against the firm, not for it.

3. Vendor due diligence files. For each third-party AI tool: written answers on what the vendor does with firm and client data, where it is stored, how long it is retained, whether it trains the vendor's models, what security representations the vendor has made, and what the vendor has committed to in the event of a breach. Amended Regulation S-P turns this from best practice into obligation, since AI vendors with access to customer information are service providers the adviser must oversee, and examiners are engaging firms on exactly this point (SEC, 2024b; SEC, 2025c). A vendor that cannot answer the diligence questions is itself a finding: document it and act on it before an examiner documents it for you.

4. Supervision records that prove human review. FINRA's 2026 report frames the expectation as human-in-the-loop validation of model outputs, with logs of prompts, outputs, and model versions (FINRA, 2026). The standard is not a policy that says recommendations receive human review. It is a record showing which outputs a named reviewer examined, when, and with what result. The December 2025 risk alert shows how staff treat the alternative: a written policy with no implementation behind it is a deficiency wearing a policy's clothing (SEC, 2025d).

5. Training records. Dates, content covered, and completion rosters for AI policy training. The distinction examiners draw is between a program that exists and a program that operates. "The policy was adopted" is an assertion; "the policy was adopted on this date, communicated on this date, and trained to completion by these people" is evidence (O'Connell, 2026).

6. Marketing substantiation, in both directions. Every AI claim in Form ADV, on the website, and in decks must trace to actual capability; that is the Delphia lesson, and the priorities commit staff to rechecking it (SEC, 2024a; SEC, 2025c). The reverse exposure gets less attention: a firm whose disclosures are silent on AI while AI drafts its client

communications and meeting summaries has a representations problem running the other way. The audit reviews what the firm claims and what it omits against what it does.

7. Incident response that reaches the AI stack. The Regulation S-P incident response program must cover unauthorized access to customer information held by AI vendors, with the 30-day notification clock attached (SEC, 2024b). Sweep letters also requested continuity plans for AI system failures and internal reports of incidents where AI use raised regulatory, ethical, or legal issues (ACA Group, 2023). If the firm has never run a tabletop exercise on an AI vendor breach, the program is untested in the place it is most likely to be needed.

Quantitative and private fund managers face deeper requests, descriptions of algorithmic trading signals, data providers, model documentation (ACA Group, 2023). The seven artifacts above are the floor for every RIA, including a firm whose entire AI estate is a notetaker and an email assistant.

5. The Hooks: Existing Rules Doing New Work

Each artifact attaches to authority that predates the technology, which is why the absence of an AI rule offers no shelter.

Fiduciary duty. Advice assembled with AI assistance is still the adviser's advice, owed with the same care and loyalty. The priorities make the application concrete: examiners test whether algorithmic and automated recommendations align with client investment profiles and stated strategies (SEC, 2025c). An adviser cannot delegate the duty of care to a model, and the exam checks that it has not.

The Compliance Rule. Rule 206(4)-7 has required reasonably designed policies and procedures, and an annual review of their effectiveness, since 2003. A compliance program with nothing to say about a technology in use at 63 percent of firms is difficult to defend as reasonably designed, and the staff has long taken the position that advisers should consider interim reviews in response to regulatory developments, a point it repeated in the December 2025 risk alert (SEC, 2025d). The annual review is the natural home for the AI program's yearly re-examination.

The Marketing Rule. Rule 206(4)-1 supplied the basis for the first AI enforcement actions and governs the testimonial, endorsement, and third-party rating mechanics that AI now generates and distributes at volume (SEC, 2024a; SEC, 2025d). Substantiation is the operative word: claims about AI capability are advertisements like any other.

Regulation S-P, with Regulation S-ID behind it. This is the one place 2026 added obligation rather than attention: incident response, breach notification, service provider oversight, and five years of records evidencing compliance, now binding on every registered adviser (SEC, 2024b). Regulation S-ID's identity theft program sits alongside it, and its red flags grow more demanding as synthetic media improves.

Books and records. Rule 204-2 requires the records that substantiate the rest, including the documentation behind testimonials and performance claims. AI outputs that inform advice, meeting summaries among them, are best treated as firm records: retained, reviewable, and held to the same accuracy standards, a posture compliance consultancies now recommend while the staff's treatment of AI-generated records develops (Hull, 2026).

One more feature of the priorities deserves notice. The Division continues to put never-examined and recently registered advisers at the front of the queue (SEC, 2025c). A firm registered in the last three years should expect its first examination to include the AI requests, with no grace period for a program still in drafting.

6. Shadow AI: The Inventory You Do Not Have

Every artifact in Section 4 presumes the firm knows what is running. The data says otherwise. Schwab found adoption proceeding "most often through individual experimentation rather than firm-wide systems" (Charles Schwab, 2026). In

the compliance testing survey, 18 percent of firms concede outright that AI tools are used informally by employees, and self-reported figures of that kind run low (IAA et al., 2025). O'Connell is blunter: he reports never having assessed an RIA where staff were not using tools the firm had never reviewed (O'Connell, 2026).

The mechanics are mundane. Consumer AI tools are free, capable, and resident on every personal phone. An advisor preparing for a client meeting pastes the household's holdings into a chatbot to draft talking points. The data leaves the firm without a trace, into a system that retains it under terms nobody read, and the firm's carefully drafted policy never touches the transaction.

Shadow AI is the worst category of exposure because it defeats the controls rather than testing them. It is invisible to the inventory, ungoverned by the policy, absent from the vendor files, and unreviewable by supervision. It also converts a clean examination into a Regulation S-P problem the firm does not know it has, since client information held in an unvetted consumer system is exactly what the incident response program exists to address.

The response has two parts. First, surface it without punishment: an amnesty survey, asking every member of staff what tools they actually use and for what, produces a truthful inventory in a way a compliance attestation never will. Second, give the demand somewhere sanctioned to go. Staff use these tools because they work. A policy that prohibits without providing pushes use further underground; an approved alternative with a fast approval lane for new requests removes the incentive to hide.

You cannot govern what you have promised to punish people for telling you about.

7. Oversight That Survives Contact With Practice

Assembling the seven artifacts is a project. Keeping them true is a design problem, and it is behavioral, not technical. The failure mode regulators keep documenting is not the absent policy but the inert one: the December 2025 risk alert is, at bottom, a catalogue of firms whose paper and practice had come apart (SEC, 2025d). My earlier work on agentic AI governance in private capital found the same pattern in a different industry: oversight systems fail at the human layer first (Coney, 2026).

Three mechanisms do most of the damage. *Automation complacency:* reviewers approve machine output because it is usually right, and the review decays into a click. Decades of human-factors research on high-automation environments document the pattern; it arrives faster when the reviewer is busy, and advisory firms are staffed to be busy. *Attestation fatigue:* quarterly certifications signed without reading, which produce records of compliance and none of the underlying behavior. *Review theater:* a "human in the loop" checkbox that no examiner can distinguish from no review at all, because nothing in the record shows what the human saw or changed.

Five design choices counter them.

Make review leave a trace. The unit of evidence is the annotation: what the reviewer corrected, queried, or rejected. A reviewer whose log shows months of approvals and not one edit is a signal to audit the review, not to celebrate the model.

Sample deeply rather than attest broadly. Blanket certification of every AI-assisted output invites fatigue and produces theater. A documented monthly sample, examined hard, with findings recorded and acted on, is stronger evidence and less work. It mirrors how firms already test marketing review and personal trading.

Name owners. A control owned by "compliance" is owned by nobody. One named person approves tools; one named person owns the inventory; reviewers are identified by name in the supervision log. Examiners ask who is responsible, and "the committee" is a poor answer.

Budget the friction. Every control adds steps, and staff route around steps. If the approved tool is slower than a consumer chatbot in a browser tab, the browser wins, and the program fails without anyone ever deciding to abandon it. The test of an acceptable use policy is whether the compliant path is also the convenient one.

Treat deviations as information. A deviation log that triggers reflexive discipline teaches concealment within a quarter. A log that triggers inquiry, why was the approved tool not used, what was missing, teaches the firm where its program does not fit the work.

Firms wanting an external scaffold for this can borrow the govern, map, measure, and manage functions of the NIST Artificial Intelligence Risk Management Framework (AI RMF 1.0), which is sector-neutral and scales down to small organizations without consultants attached (NIST, 2023). The structure matters less than the property it protects: an oversight system that still produces honest records in month eighteen, after the novelty has worn off.

8. The 90-Day Build

None of this requires a year or a seven-figure budget. It requires sequence. The build below fits a mid-sized RIA in a quarter; a solo practice compresses it into weeks, because the artifacts scale down gracefully. What does not scale down is their absence.

Days 1 to 30: find everything, then draw the perimeter. Run the amnesty survey across every member of staff, including operations and any affiliates. Census the embedded AI in existing vendor platforms, the CRM, the meeting tool, the email stack, since vendors have been switching these features on by default. Classify every discovered tool by the data it touches. Draft the acceptable use policy against the real inventory rather than a template, and stand up the approval lane with a named owner and a committed turnaround. Freeze new tools until the lane exists.

Days 31 to 60: vendors and supervision. Build diligence files for the highest-risk vendors first, those touching client personal information. Send the diligence questionnaire; chase written answers on data handling, storage, retention, model training, and breach notification; escalate the vendors that will not answer. Map each relationship to the Regulation S-P obligations and update the incident response program so it names AI systems explicitly. Rewrite the supervision procedure to specify which AI-assisted outputs get reviewed, by whom, on what sampling basis, and where the record lives. Audit every AI claim, and every material AI silence, in Form ADV and marketing materials against the inventory.

Days 61 to 90: make it operate, then prove it operates. Train all staff on the policy and document completion by name and date. Run the first sampled supervision cycle and record findings, including the uncomfortable ones; a first cycle with zero findings means the sample was too polite. Tabletop one scenario: an AI notetaking vendor reports unauthorized access to meeting transcripts containing client information, and the firm walks the 30-day notification clock end to end. Brief leadership on results, calendar the AI item into the annual Rule 206(4)-7 review, and set the quarterly cadence for tool approvals and policy updates.

At day 90 the firm holds the seven artifacts, each with a date, an owner, and a first cycle of operating evidence behind it. That is what an examiner means by a program.

9. Conclusion and Recommendations

The deregulatory turn at the Commission was real, and it changed the form of AI oversight rather than the fact of it. Regulation arrived by examination instead of rulemaking: an enforcement precedent in 2024, an amended privacy rule with teeth, named priorities in 2025, and a final compliance deadline that passed on June 3, 2026. The firms still waiting for an AI rule to tell them what good looks like have the sequence backwards. The examiner is already asking.

The recommendations reduce to six.

1. Build the inventory first, through amnesty rather than attestation, and refresh it quarterly. Every other control inherits its accuracy.

2. Write the acceptable use policy against your actual tool population, name an approver, and commit to a turnaround fast enough that staff use the lane instead of routing around it.

3. Map every AI vendor to the amended Regulation S-P obligations: diligence files, breach notification terms, and an incident response program that names the AI stack.

4. Convert supervision from assertion to record. Sampled review with annotations and named reviewers; no review that leaves no trace.

5. Audit AI representations in both directions, the claims and the silences, against practice, and recheck at every marketing refresh and ADV update.

6. Design for the human failure modes: traceable review, deep sampling over broad attestation, named owners, a friction budget, and deviation logs that produce inquiry rather than concealment.

There is a commercial argument hiding inside the compliance one. Institutional clients and sophisticated prospects have started asking advisers how they govern AI, and firms with a documented answer are winning mandates from firms without one (O'Connell, 2026). The binder you assemble for the examiner is the same one that closes the diligence question. Build it once, run it honestly, and the exam becomes what Chairman Atkins says it should be: a constructive conversation, with records on your side of the table.

References

ACA Group. (2023). *SEC conducts sweep of AI use by investment advisers.* Compliance alert, 14 December 2023. ACA Group.

Charles Schwab. (2026). *Schwab study reveals RIA AI adoption more than doubles, but most firms still in early stages.* Press release, 22 January 2026. Charles Schwab & Co., Inc.

Coney, L. (2026). *Agentic AI governance in private equity: A behavioral framework for autonomous decision systems.* WorkWise Solutions, Q1 2026.

FINRA. (2024). *Regulatory Notice 24-09: FINRA reminds members of regulatory obligations when using generative artificial intelligence and large language models.* Financial Industry Regulatory Authority, 27 June 2024.

FINRA. (2026). GenAI: Continuing and emerging trends. In *2026 FINRA Annual Regulatory Oversight Report.* Financial Industry Regulatory Authority.

Hull, S. (2026). *AI compliance for wealth management firms and RIAs: What to know in 2026.* Ncontracts Nsight blog, 31 March 2026.

Investment Adviser Association, ACA Group, & Yuter Compliance Consulting. (2025). *2025 Investment Management Compliance Testing Survey.* July 2025.

National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* NIST AI 100-1, January 2023.

O'Connell, J. (2026). No new rules doesn't mean no examination. *WealthManagement.com*, 8 May 2026.

SEC. (2023). *SEC proposes new requirements to address risks to investors from conflicts of interest associated with the use of predictive data analytics by broker-dealers and investment advisers.* Press release 2023-140, 26 July 2023. U.S. Securities and Exchange Commission.

SEC. (2024a). *SEC charges two investment advisers with making false and misleading statements about their use of artificial intelligence.* Press release 2024-36, 18 March 2024. U.S. Securities and Exchange Commission.

SEC. (2024b). *Regulation S-P: Privacy of consumer financial information and safeguarding customer information.* Final rule, Release No. 34-100155, 16 May 2024. U.S. Securities and Exchange Commission.

SEC. (2025a). *Withdrawal of proposed regulatory actions.* Release No. 33-11377, 12 June 2025. U.S. Securities and Exchange Commission.

SEC. (2025b). *SEC creates task force to tap artificial intelligence for enhanced innovation and efficiency across the agency.* Press release 2025-103, 1 August 2025. U.S. Securities and Exchange Commission.

SEC. (2025c). *Examination priorities: Fiscal year 2026.* Division of Examinations, 17 November 2025. U.S. Securities and Exchange Commission.

SEC. (2025d). *Risk alert: Additional observations regarding advisers' compliance with the Advisers Act Marketing Rule.* Division of Examinations, 16 December 2025. U.S. Securities and Exchange Commission.

About the Author

Dr. Leigh Coney is the Founder and Principal Consultant of WorkWise Solutions. His doctoral research examined how people interact with emerging technology, and he has spent over a decade at the intersection of AI, behavioural science, and organisational design. His work focuses on decision making in high stakes environments, with particular attention to why sophisticated AI systems fail to achieve adoption and how governance and deployment systems can be designed to account for human cognitive and social dynamics.

Dr. Coney's work is distinguished by its integration of behavioural science with practical technology deployment. He advises registered investment advisers, private equity firms, family offices, private credit funds, independent sponsors, and investment banks on AI strategy, governance design, and psychology informed change management.

This paper is part of an ongoing research series on responsible AI adoption in financial services. Previous publications include *The Data Prerequisite: Minimum Viable Data Infrastructure for AI in Portfolio Companies* (Q2 2026), *The AI Exit Premium: Preparing Portfolio Companies for Buyer Scrutiny* (Q2 2026), *Diligencing the AI Story: Separating Capability from Theater in Target Companies* (Q2 2026), and *Agentic AI Governance in Private Equity* (Q1 2026).

About WorkWise Solutions

WorkWise Solutions builds secure, purpose built AI systems for registered investment advisers, wealth managers, private equity, venture capital, family offices, private credit funds, and investment banking firms. The firm specialises in zero retention AI architecture that ensures proprietary client and portfolio data never train public models.

WorkWise's approach is grounded in a core insight: most AI implementations fail because of broken workflows and poor adoption strategies, not because of the technology itself. Every engagement integrates behavioural science and organisational psychology into the technical design, so that AI systems become a durable part of how investment teams actually work.

Contact: leigh@workwisesolutions.org