

Artificial Intelligence Use Policy & Procedures

The same sample minimum framework — marked up by the compliance seat that has to live with it. Every red box prices a provision: the trap an examiner will find, the homework it commits you to by hand, and what dissolves under a governed AI layer.

Adopted by: [FIRM NAME] • Effective date: [DATE] • Policy owner: Chief Compliance Officer

Version: 1.0 (sample) • Next review due: [DATE + 1 YEAR]

READ THIS FIRST

This document is a free educational sample published by RIAGuardrAils, a technology vendor. It is marketing material, not legal, compliance, or investment advice, and adopting it does not create compliance with any rule. Every firm's obligations depend on its business, and this framework must be tailored and reviewed by your compliance counsel before adoption. No attorney-client relationship is created. One more thing, stated plainly in Section 11: a policy you cannot evidence is a liability, not a defense — before adopting any provision, decide how you will produce the records it promises.

1. Purpose and scope

- 1.1** This policy governs the use of artificial intelligence systems — including generative AI and AI-enabled features embedded in other software — by all supervised persons of [Firm Name] (the "Firm") in connection with Firm business.
- 1.2** This policy forms part of the Firm's compliance program adopted under Rule 206(4)-7 of the Investment Advisers Act and is subject to the Firm's annual review.
- 1.3** Nothing in this policy authorizes any use of AI that conflicts with the Firm's fiduciary duty, Form ADV disclosures, client agreements, or applicable law.

2. Definitions

- 2.1** "AI System" means any software that generates content, predictions, recommendations, summaries, or decisions from data — including chat assistants, meeting note-takers, transcription tools, drafting aids, analytics copilots, and AI features inside CRM, portfolio, or planning software.
- 2.2** "Approved AI Tool" means an AI System listed on the Firm's Approved AI Tool Inventory (Appendix A) after the diligence in Section 10.
- 2.3** "Client Information" means any nonpublic personal information of a client or prospect as defined under Regulation S-P, and any account, holding, transaction, or planning data attributable to an identifiable client.

3. Governance and accountability

- 3.1** The Chief Compliance Officer (or designee) administers this policy, maintains the Approved AI Tool Inventory, and is the approval authority for any new AI use.

- 3.2** The Firm maintains a current inventory of every AI System in use for Firm business, including embedded AI features of existing vendors. Unknown tools are treated as unapproved.
- 3.3** Exceptions to this policy require the CCO's prior written approval and are logged with the rationale and any conditions.

EVIDENCE THE EXAM WILL ASK FOR

- ▶ The current Approved AI Tool Inventory, with approval dates and the diligence file behind each entry.
- ▶ The exception log — or an affirmative record that no exceptions were granted.

4. Approved tools and permitted use

- 4.1** Supervised persons may use only Approved AI Tools for Firm business, and only through Firm-provisioned accounts. Personal accounts on consumer AI services may not be used for any Firm purpose.
- 4.2** Permitted uses include drafting and editing non-client-specific content, research and education, summarizing Firm-approved documents, and any use expressly approved for a specific tool in Appendix A.
- 4.3** All AI output used in Firm business is reviewed by a qualified person before use. AI does not render investment advice to clients; every recommendation remains the responsibility of the advisor of record.

EVIDENCE THE EXAM WILL ASK FOR

- ▶ Access records showing Firm-provisioned accounts — and how the Firm would detect use of personal accounts.
- ▶ For a sampled AI-assisted deliverable: who reviewed it before it was used, and when.

FROM THE COMPLIANCE SEAT — ANNOTATED EDITION

⚠ THE TRAP "Firm-provisioned accounts" assumes consumer AI tools offer firm provisioning, role scoping, and admin visibility. Most don't — which makes 4.1 a rule about tools that can't implement it, enforced by hoping.

THE HOMEWORK You become the IT administrator, license manager, and access auditor of a pile of point AI subscriptions — each with its own account model, none with your org chart.

✓ **UNDER A GOVERNED LAYER** One platform provisions every seat by role. Advisors, executives, and compliance each get a scoped seat — and "how would you detect a personal account" becomes "there's nothing useful outside the governed door."

5. Prohibited uses

- 5.1** Entering Client Information into any AI System that is not an Approved AI Tool expressly authorized for Client Information. This prohibition covers pasting, uploading, dictating, and screen-sharing.
- 5.2** Using AI to generate advertisements, performance claims, or testimonials except through the Firm's marketing-review process under the Marketing Rule.
- 5.3** Representing AI capabilities the Firm does not actually use — in marketing, RFPs, or client communications — or overstating the role of AI in the Firm's processes ("AI-washing").
- 5.4** Using AI Systems to make, or appear to make, suitability or best-interest determinations without documented human judgment.

6. Accuracy of AI output; use in advice

- 6.1** Generative AI can produce plausible but false output ("hallucination"), including invented facts, figures, citations, and account details. Supervised persons are trained to treat all generative output as unverified draft.
- 6.2** No AI-generated number, holding, transaction, performance figure, tax figure, or client-specific fact may be used in Firm business — including client communications, reviews, and reports — until it has been verified against the Firm's systems of record. The verification is documented for client-facing deliverables.
- 6.3** AI output may inform, but never constitutes, investment advice. The advisor of record independently forms and documents the basis for every recommendation.
- 6.4** The Firm's Form ADV and client-facing descriptions of its AI use are kept accurate and current, including material changes in how AI touches client accounts or communications.

EVIDENCE THE EXAM WILL ASK FOR

- For a sampled client deliverable containing figures: the verification trail back to the system of record, and who performed it.
- Training records showing supervised persons were instructed on hallucination risk and the verification requirement.

FROM THE COMPLIANCE SEAT — ANNOTATED EDITION

⚠ THE TRAP 6.2 is the sentence that quietly costs the most. Adopt it and every number in every AI-assisted draft — every balance, every return, every RMD — must be traced to a source system by a human, every time, forever. Skip the verification once, ship one hallucinated figure to a client, and 6.2 is the provision you're charged under.

THE HOMEWORK A per-deliverable verification workflow, a documentation habit nobody keeps under deadline, and a standing bet that no advisor ever trusts a fluent paragraph with a wrong number in it.

✓ **UNDER A GOVERNED LAYER** The model never computes a figure. Numbers come only from governed tools reading the systems of record, provenance-tagged — an unverified figure structurally cannot appear on the screen, so 6.2 is satisfied by architecture, not willpower.

7. Client information and privacy

- 7.1** Client Information is safeguarded under the Firm's Regulation S-P policies. AI Systems that receive Client Information must sit inside the Firm's safeguards: access-controlled, encrypted in transit and at rest, and covered by the Firm's incident-response program.
- 7.2** Where technically supported, client identifiers are minimized or masked before any AI System processes them (for example, account numbers surfaced only as last-4).
- 7.3** Each supervised person's AI access is limited to the Client Information appropriate to their role. Access rights are reviewed at least annually and upon role change.

EVIDENCE THE EXAM WILL ASK FOR

- For each tool authorized for Client Information: where the data resides, the vendor agreement, and the safeguard review.
- The access-scoping model — and a test showing a supervised person cannot retrieve client data outside their role.

FROM THE COMPLIANCE SEAT — ANNOTATED EDITION

⚠ THE TRAP "Where technically supported" is doing heroic work in 7.2 — consumer tools don't mask, and 7.3

promises role-scoping that chat apps have no concept of. You've written provisions your tool inventory cannot execute.

THE HOMEWORK *Manual redaction habits, per-tool configuration reviews, and an annual access review over systems with no access model to review.*

✓ **UNDER A GOVERNED LAYER** *Masking to last-4 is the default path, not an aspiration; seats are scoped fail-closed and the scoping is testable on demand — 7.2 and 7.3 describe the architecture, not a hope.*

8. Books and records

- 8.1** Records of AI use that relate to the Firm's advisory business — including AI-assisted client communications, AI-generated content used in Firm business, and records of AI access to Client Information — are preserved under Advisers Act Rule 204-2 for the required periods.
- 8.2** AI interactions conducted through channels the Firm cannot capture and produce (personal accounts, unapproved apps, personal devices) are prohibited precisely because they create records the Firm cannot maintain.
- 8.3** The Firm can identify, retrieve, and produce its AI-use records upon regulatory request within a reasonable time.

EVIDENCE THE EXAM WILL ASK FOR

- ▶ A produced sample: the record of AI interactions touching client data for a chosen period — who asked, what was accessed, when, and the outcome.
- ▶ Where those records live, the retention schedule, and who can produce them.

FROM THE COMPLIANCE SEAT — ANNOTATED EDITION

⚠ **THE TRAP** *8.3 promises production of records that, in a consumer-tool program, live in personal accounts you don't control. You have adopted a books-and-records obligation over books you don't hold.*

THE HOMEWORK *Chasing exports from a dozen apps, or admitting the record doesn't exist — under Rule 204-2, the record you can't produce is the finding.*

✓ **UNDER A GOVERNED LAYER** *Every AI data request writes an audit entry as it happens — who asked, what was accessed, when, the outcome. 8.3 is answered by a query against a ledger the Firm owns.*

9. Marketing, communications, and AI claims

- 9.1** Statements about the Firm's use of AI — in Form ADV, marketing, RFPs, and client communications — are accurate, substantiated, and reviewed under the Firm's advertising procedures.
- 9.2** AI-generated marketing content follows the same review-and-approval path as any other advertisement under the Marketing Rule.

10. Vendor due diligence and oversight

- 10.1** Before approval, each AI vendor is assessed for: data usage and model-training rights over Firm data; data residency and segregation; security posture; retention and deletion terms; subprocessors; and exit rights, including return or destruction of Firm data on termination.
- 10.2** AI vendors receiving Client Information are treated as service providers under the Firm's amended Regulation S-P program, including oversight and incident-notification obligations.

- 10.3** Approved vendors are re-reviewed on a defined cycle and upon material change. Every retired AI subscription is removed from the inventory and offboarded — credentials revoked, data disposition confirmed.

EVIDENCE THE EXAM WILL ASK FOR

- ▶ The diligence file for each vendor authorized for Client Information — including the data-training and exit-rights answers.
- ▶ Proof of offboarding for any retired AI vendor: access revoked, data disposition confirmed in writing.

11. Monitoring, testing, and evidence

- 11.1** The CCO (or designee) periodically reviews the Firm's AI use for conformity with this policy, on a documented cadence, using records the Firm actually possesses.
- 11.2** Monitoring includes, at minimum: sampling AI interactions involving Client Information; testing that access scoping works as described; reviewing the exception log; and confirming the tool inventory is complete.
- 11.3** Violations are documented, escalated, and remediated under the Firm's compliance procedures. The absence of any detected violations is itself recorded with the review that reached it.
- 11.4** The Firm acknowledges the operative standard: a written policy without produceable monitoring records is evidence of a compliance program that exists on paper only. Each provision of this policy is adopted together with the record that proves it.

EVIDENCE THE EXAM WILL ASK FOR

- ▶ The monitoring calendar and the completed reviews — dates, sample, findings, and the reviewer.
- ▶ The record demonstrating each review actually examined AI-use data, not attestations alone.

FROM THE COMPLIANCE SEAT — ANNOTATED EDITION

⚠ THE TRAP 11.2 orders you to sample interactions you cannot see and test scoping that doesn't exist. When the examiner asks "how do you know?", an attestation is what advisors say — and the certification you can't sign yourself becomes the finding. That's the off-channel sequence, next channel.

THE HOMEWORK A monitoring calendar over an empty data set, or device-imaging your own advisors — the two options the sweep firms faced.

✓ **UNDER A GOVERNED LAYER** Monitoring is a query: "show me every AI data request last quarter — who asked, what was touched, what came back." The completed review becomes a standing record, produced on demand.

12. Training and attestations

- 12.1** Supervised persons receive training on this policy at adoption, at hire, and at least annually thereafter.
- 12.2** Each supervised person attests annually to compliance, including that they have not entered Client Information into unapproved AI tools.
- 12.3** Attestations supplement, and do not substitute for, the monitoring records described in Section 11. The Firm does not rely on self-certification as its sole evidence that this policy is followed — an acknowledgment answers what supervised persons say; monitoring answers how the Firm knows.

EVIDENCE THE EXAM WILL ASK FOR

- ▶ Completed attestations for every supervised person — paired with the Section 11 monitoring record that corroborates them.

13. Incident response

13.1 Suspected policy violations, AI-related data exposure, or vendor incidents are reported to the CCO immediately and handled under the Firm's Regulation S-P incident-response program, including any required client notification.

14. Annual review

14.1 This policy, the tool inventory, the monitoring record, and the training and attestation files are reviewed as part of the Firm's annual review under Rule 206(4)-7, and the review is documented.

EVIDENCE THE EXAM WILL ASK FOR

- The annual-review workpaper section covering AI — what was examined, what changed, and who signed.

Appendix A. Approved AI tool inventory (template)

Tool / vendor	Approved uses	Client info authorized?	Diligence file date	Approved by / date

Appendix B. The evidence checklist — one page for the exam file

If the Firm permits AI use, expect to be asked to produce:

- The written AI policy, with adoption date and the annual-review trail.
- The complete AI tool inventory, including embedded AI features of existing vendors.
- Vendor diligence files for every tool touching Client Information — data-training rights, residency, exit terms.
- The record of AI interactions with Client Information: who asked, what was accessed, when, and the outcome.
- The access-scoping model and a demonstration that it fails closed.
- The verification trail for AI-generated figures used in client deliverables.
- Monitoring reviews on a documented cadence, showing real data was examined.
- The exception log and the violation log — or the affirmative record that each is empty.
- Training materials and annual attestations for every supervised person.
- Incident-response records for any AI-related event, including Reg S-P analysis.

Ten artifacts. Under a manual program, each is a project. Under a governed AI layer, items 4 through 7 are byproducts of the architecture — the system produces its own evidence. That difference is the entire reason this sample is free.

Adoption

Adopted for [Firm Name] by:

Chief Compliance Officer

Date

© 2026 RIAGuardrAlls · riaguardrails.com · This annotated edition may be shared in unmodified form with attribution. Your AI. Your Data. Your Rules.