

Helmsted

THE RIA'S CHARTBOOK · NO. 01


FOR INDEPENDENT RIAs · REGULATION S-P

Client Data and the Law.

What Regulation S-P expects of independent RIAs as AI moves into their workflows, and how to be ready before an incident.

INSIDE THIS GUIDE

- 01 AI is moving from answers to action
- 02 Follow the client information through the work
- 03 What Reg S-P brings into focus
- 04 Personal AI accounts are not a safeguard
- 05 Prepare for the incident before the notice arrives
- 06 Treat incoming documents as information, not authority
- 07 Start with one workflow your firm can defend



SECTION 01

AI is moving from answers to action

At independent RIAs, AI is moving from answering questions to doing work: preparing meetings, updating records, coordinating follow-up, and carrying out the steps a client has approved. The adviser and client decide what should happen. AI increasingly carries the authorized steps forward.

That shift changes where client information goes. A tool that drafts one email touches a little data once. A system that prepares, maintains, and coordinates work touches more records, more often, and passes them through more services. Every one of those paths is customer information the firm is responsible for protecting.

For an advisory firm, the useful question becomes: **Can this system carry out the work we intend, with the information and authority we intend to give it, and can we show how that information is protected?**

Before a firm can delegate the work, it must understand and protect the information that makes the work possible. That is the subject of this installment. How authority grows from there is the subject of No. 03.

The adviser and client decide. AI carries the authorized work forward.

This guide focuses on the rules and the readiness behind that answer. It brings selected Regulation S-P requirements together with a practical view of AI workflows. The workflow recommendations are editorial guidance, not a claim that the regulation prescribes a particular AI architecture.

SECTION 02

Follow the client information through the work


A tax return is an easy place to begin. Imagine an adviser using an AI application to prepare for a retirement-planning meeting with the Morgans, a fictional couple three years from retirement who appear throughout this series. Depending on the implementation, the file could pass through document extraction, application storage, a model provider, a search index, and an integration that adds information to the firm's CRM.

The result can create more records: extracted facts, planning notes, a meeting agenda, a draft message, and an activity history. Protecting the original upload is only one part of understanding the workflow.


Draw the path from receipt to final destination. At each step, identify the information involved, its purpose, who can access it, and what happens to it afterward. Include support access, operational logs, and backups where applicable. An AI-generated summary can still contain identifying financial information even when the original document is no longer attached.

The map should reflect the configuration the firm actually uses. Adding a meeting recorder, enabling a new connector, or allowing external messages can change both the recipients of information and the consequences of an error.

For CFP® professionals, understanding how information moves through the technology platform is also part of CFP Board's technology guidance. Its guide connects knowledge of the platform with appropriate evaluation of the results it produces. [1]



The adviser and client decide what should happen. AI can carry the authorized work forward.



PART II · THE RULES

The rules already reach your AI tools.

SECTION 03

What Reg S-P brings into focus

The 2024 amendments to Regulation S-P cover SEC-registered investment advisers and other specified institutions. A state-registered adviser's obligations require a separate analysis; this guide does not treat all advisers as subject to the same SEC rule. The compliance dates were December 3, 2025 for larger entities and June 3, 2026 for smaller entities. Both have passed. [2]

For covered firms, key requirements include written incident-response procedures, service-provider oversight, customer notification when required, safeguards and disposal controls, and compliance records. The rule reaches customer information handled on the firm's behalf. [2]

Examiners have said where they will look. The SEC Division of Examinations' fiscal 2026 priorities state that it will examine firms for compliance with the amended rule, and that "examinations will focus on firms' policies and procedures, internal controls, oversight of third-party vendors, and governance practices." The same priorities say the Division will review registrants' representations about their AI capabilities for accuracy, and assess whether firms have adequate policies and procedures to monitor or supervise their use of AI technologies. [3]

For a firm adopting AI, the practical implication is to connect the policy with an actual workflow. Someone should be able to show which provider was evaluated, which use was approved, what has changed since approval, and how an incident would reach the right people.

SECTION 04

Personal AI accounts are not a safeguard

Reg S-P's safeguards and service-provider oversight do not stop at the firm's own systems. When customer information goes to an AI model provider, that provider is handling it on the firm's behalf, and the arrangement has to hold up to the same questions as any other vendor.

Consumer AI accounts were not built for that. Terms for individual plans are written for individuals, and they can differ sharply from business terms. One major provider, for example, states that content from its individual chat services may be used to train its models unless the user opts out, while inputs and outputs from its business, enterprise, and API offerings are not used for training by default. [4] Its business and enterprise tiers add workspace administration and retention controls the firm can set.

[5]

The regulation does not name a plan. But a firm would struggle to show reasonably designed safeguards and meaningful provider oversight for client information typed into personal accounts it does not control, under terms written for consumers.

Our recommendation, beyond what the rule states: keep client information in arrangements the firm has evaluated and can govern, such as business or enterprise agreements or API access under appropriate terms, and oversee that provider like any other. No. 02 covers how to evaluate the evidence.

SECTION 05

Prepare for the incident before the notice arrives


Suppose an AI service provider reports unauthorized access to a system used in the firm's document workflow. The provider is still investigating which records were involved.

The firm needs a practiced route from that message to action. It should know who coordinates the response, which technical and legal resources to engage, where the relevant records are, and who can pause affected workflows. Under the final rule, receipt of the specified service-provider notice requires the covered institution to initiate its incident-response program. [6, section II.A.4]

Two notification clocks deserve particular care. They have different triggers and different recipients.

The customer-notice rule includes an exception based on a reasonable investigation finding no actual or reasonably likely use of sensitive customer information causing substantial harm or inconvenience, and limited permitted delays. [6, sections II.A.3–4]

The firm should not add the two outer limits together or wait for a vendor deadline before responding to facts it already knows. [6, section II.A.4]



NIST's incident-response guidance frames the work as preparation, response, and recovery, and is a useful supporting reference for the exercise below. It is voluntary guidance, not a Reg S-P requirement. [9] Run the fictional incident as a short exercise, using the Morgans' file as the affected record. Uncertainty should be recorded and investigated. "The provider does not train on our data" does not resolve an access incident. "We have not seen misuse" is a fact to evaluate, not a complete notice determination.

Four questions structure the exercise, and they tend to surface gaps before a real event: an outdated contact, an uninventoried integration, a log no one can retrieve.

EXERCISE QUESTION	WHAT THE TEAM SHOULD PRODUCE
What happened?	The provider's notice, the affected system, and when the firm knew.
Who owns the response?	The coordinator, who can pause work, and who owns the notice analysis.
What evidence is needed?	The information map, access logs, and the households involved.
What can be contained now?	Paused connectors and revoked access, each recorded.

SECTION 06

Treat incoming documents as information, not authority

An AI workflow may read messages, uploaded files, and external pages that the firm did not write. Those materials can contain text that looks like an instruction to the system.

NIST describes indirect prompt injection as an attack in which instructions are placed in material an AI application is likely to retrieve. It identifies this as an information-security risk for generative AI systems. [7, section 2.9]

Consider a fictional example: a document contains a line telling the application to email the entire household file to a new address. Reading that line should not authorize the action. The firm's approved instructions and permissions must control what the system can do.

A practical evaluation should include misleading source material, unexpected recipients, and requests outside the task's scope. Test the full workflow, including connected tools. A response that sounds appropriately cautious is insufficient evidence if the application can still take the prohibited action.

Useful controls to evaluate include limiting tool permissions, verifying destinations, separating household access, and requiring escalation when an action exceeds its authority. These are implementation recommendations, not a promise that any individual control eliminates the risk.


Start with one workflow your firm can defend

The steps in this section are Helmsted's recommended operating practice, not requirements stated in the rule. Choose one recurring AI-enabled process, such as meeting preparation or collecting missing planning documents, and name the person accountable for it.

Map the information involved and every service that receives it. Record which provider was evaluated for each step, which use was approved, and what evidence the firm holds. Write down the incident path, then run the four-question exercise above with a fictional household and provider notice. Keep the results with the firm's compliance records, and revisit them whenever the workflow, a connector, or a provider changes.

KEEP IN THE WORKFLOW FILE	WHAT IT SHOWS AN EXAMINER
Accountable owner	Who answers for the workflow.
Information map	Where customer information goes, including logs, indexes, and backups.
Providers and approvals	Which services were evaluated, for what use, with what evidence.
Incident path	Who is notified, who can pause the work, and who owns the notice analysis.
Exercise results	That the response was tested, and what was fixed.
Change log	When the workflow, a connector, or a provider changed, and what was re-reviewed.

One workflow is a starting point, not a ceiling. Each workflow the firm can map, test, and defend is one it can trust with more authority, and the next one is easier to add.



IN CLOSING

From policy to practice.

The rules ask a firm to know where customer information goes, to oversee the providers that handle it, and to respond in order when something goes wrong. Every additional tool multiplies that work. As AI carries more of the firm's work, these questions become the operating rhythm of a firm that delegates.

This is where Helmsted's structure matters. Helmsted brings the technology layer into one secure platform built for the independent RIA: the household record, the permissions a firm grants, and the record of what was done live in one place, with fewer paths to trace when it counts. Our subprocessors are listed publicly at helmsted.ai/security/subprocessors, and our Data Processing Addendum addresses safeguards, incident response, and AI-provider arrangements with no model training and zero retention of customer inputs and outputs.

Our security program pairs continuous monitoring through Vanta with an independent SOC 2 Type I examination by Sensiba LLP, which issued a clean report covering Security and Confidentiality, available under NDA.

Know where client information goes. Be ready before the notice arrives.

**See how Helmsted keeps the household,
and its record, in one place.**

Thirty minutes with the team, and every security question answered.

[Book a demo · helmsted.ai/demo](https://helmsted.ai/demo)

Sources and further reading

Research checked September 24, 2026. Bracketed numbers in the text refer to the sources below.

1. **CFP Board, Guide to CFP Board's Technology Standard (2024)**. Relevant passages: printed pp. 2–5 on AI, professional judgment, and the technology platform; pp. 7–10 on diligence and evaluation of outcomes. Professional guidance for CFP® professionals; not an SEC rule. [Read the guide](#).
2. **SEC, Enhancements to Regulation S-P: A Small Entity Compliance Guide (2024)**. Relevant headings: who is affected; scope of information; requirements; compliance dates. A staff summary, not a substitute for the rule. [Read the guide](#).
3. **SEC Division of Examinations, Fiscal Year 2026 Examination Priorities (November 17, 2025)**. Sections on Regulation S-P and S-ID, cybersecurity, and emerging financial technology, including AI. Staff priorities, not rules. [Read the priorities](#).
4. **OpenAI, How your data is used to improve model performance (Help Center, accessed September 2026)**. Content from individual ChatGPT services may be used to train models unless the user opts out; inputs and outputs from ChatGPT Business, Enterprise, Edu, and the API are not used by default. An example of how one provider separates consumer and business terms. [Read the article](#).
5. **OpenAI, Enterprise privacy (accessed September 2026)**. No training on business data by default for ChatGPT Business, Enterprise, and the API Platform; retention controls for Enterprise workspaces. [Read the page](#).
6. **SEC, Regulation S-P: Privacy of Consumer Financial Information and Safeguarding Customer Information, Release Nos. 34-100155, IA-6604, IC-35193 (May 16, 2024), conformed to the Federal Register version**. Relevant sections: II.A.3, customer notices; II.A.4, provider oversight, reporting, and removal of the proposed written-contract requirement; II.B, scope; II.C, records; II.F, compliance period. Adopted rule text: 17 CFR 248.30 (a) (3) – (5). [Read the adopting release and rule](#).
7. **NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1 (July 2024)**. Section 2.2 discusses confabulation and next-token prediction; section 2.9, Information Security, discusses direct and indirect prompt injection; the profile also offers voluntary risk-management actions. It does not establish a Reg S-P compliance certification. [Read the profile](#).
8. **CFP Board, Generative AI Ethics Guide: A Checklist for Upholding the Code and Standards (February 25, 2025)**. Additional reading for CFP® professionals considering AI use. [View the publication and checklist link](#).
9. **NIST, Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile, SP 800-61 Rev. 3 (April 2025)**. Organizes incident preparation, response, and recovery around the CSF 2.0 functions. Voluntary supporting guidance, not a Reg S-P requirement. [Read the publication](#).

Also in The RIA's Chartbook. No. 01, *Client Data and the Law*: what Reg S-P expects as AI enters the workflow. No. 02, *Every Vendor Is Part of Your Safeguards*: evaluating the technology that holds client information. No. 03, *Faster Check-ins, or Better Health?*: the harness behind proactive, AI-enabled advice.

Educational information, not legal advice. Applicability and incident decisions depend on the firm's circumstances and other relevant requirements. Workflow examples describe an operating vision, not a representation that all features are currently available in Helmsted. Images are AI-generated editorial illustrations; people and settings are fictional.

Helmsted

One secure, AI-native platform for the independent RIA. Learn more at helmsted.ai, or write to hello@helmsted.ai.