SurgeONE.ai
AI-Powered. Practitioner-Led

Defensible AI for RIAs

How Compliance Leaders Can Adopt AI Responsibly Without Creating Regulatory Risk

An educational guide for Chief Compliance Officers, CEOs, & CIOs at small and mid-sized Registered Investment Advisors

Introduction | AI Adoption Is No Longer Theoretical. Governance Cannot Be Either.

AI has moved quickly from experimentation to daily use inside financial firms. Employees are using AI-enabled tools to draft policies, summarize meetings, review documents, write communications, analyze data, support supervision, and improve productivity. Some of these tools are intentionally adopted by the firm. Others enter through ordinary software upgrades, vendor platforms, browser extensions, personal subscriptions, or employee workarounds.

For RIAs and broker-dealers, this creates a difficult operating reality. AI can improve efficiency, but it also introduces risk across privacy, recordkeeping, supervision, cybersecurity, vendor oversight, accuracy, bias, and regulatory accountability. The problem is not simply whether a firm uses AI. The more pressing issue is whether the firm knows where AI is being used, what data is being exposed, who is reviewing the output, and how the firm would defend its controls during an examination.

The regulatory environment remains fragmented, but that does not mean expectations are unclear. Regulators may not have a single AI rulebook, but they already have well-established rules and principles for supervision, privacy, books and records, fiduciary duty, cybersecurity, vendor oversight, and investor protection. AI will be evaluated through those existing frameworks. A firm cannot avoid accountability by saying the technology made the decision, drafted the document, created the error, or exposed the data.

www.surgeone.ai info@surgeone.ai +1 650-294-8889

As one regulatory expert put it, “The model is not responsible, and the firm and the representatives or advisors charged with overseeing the AI, they are responsible.”

This is the central premise of defensible AI. Responsible adoption does not require firms to avoid AI. It requires firms to adopt AI in a way that can be explained, supervised, documented, tested, and corrected. For RIAs, especially small and mid sized firms with limited compliance resources, the objective is not to build a theoretical AI governance framework that sits on a shelf. The objective is to create a practical operating model that allows the firm to use AI while preserving control.

That means firms need to answer several concrete questions:

- What AI tools are being used across the organization?
- Which tools are approved, restricted, or prohibited?
- What client, investor, employee, firm, or trading data can be entered into each tool?
- Who reviews AI-generated output before it is used?
- How does the firm test whether AI tools are accurate, current, biased, or hallucinating?
- How are AI use, approvals, exceptions, vendor reviews, and testing documented?
- How would the firm respond if AI caused a data breach, inaccurate recommendations, deficient procedure, or supervisory failure?

The answers to those questions form the basis of a defensible AI program.

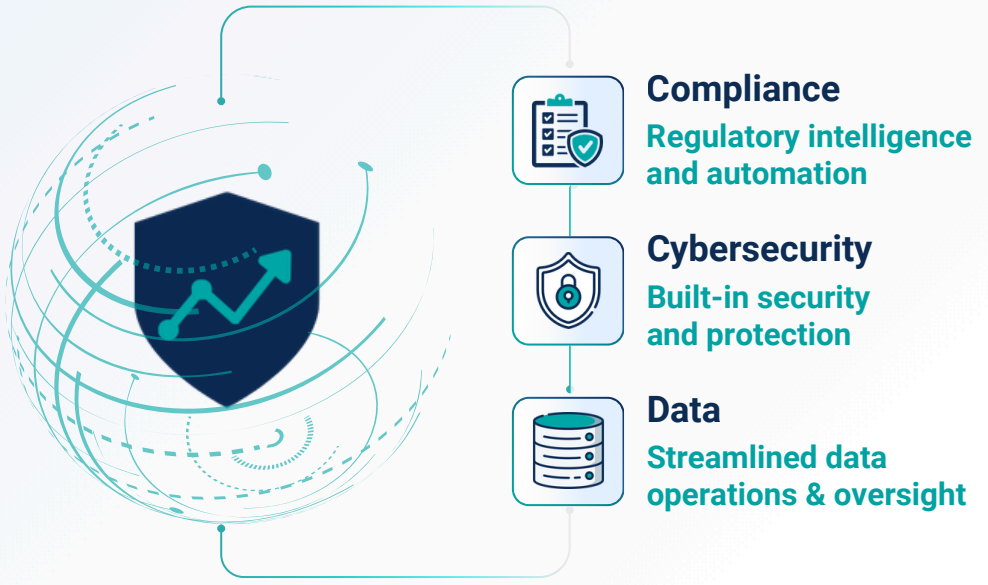

SurgeONE.ai
AI-Powered. Practitioner-Led

Defensible AI Begins with Control, Not Experimentation

AI adoption often begins informally. Someone uses ChatGPT to draft an email. A compliance team uses an AI assistant to summarize a policy. An advisor uses an AI note-taker. A vendor adds AI functionality to an existing platform. A team member uses a personal Claude, Gemini, Copilot, Grok, or ChatGPT account to speed up work.

At first, these use cases may appear low risk. But from a compliance perspective, the risk profile changes as soon as firm data, client information, supervisory processes, records, recommendations, or regulated communications are involved. The firm then needs to move from informal usage to defined governance.

A defensible AI program begins with a simple but often overlooked requirement: the firm must know where AI is being used. Without inventory, there can be no meaningful supervision. Without supervision, there can be no defensibility.

The starting point is to build a clear policy framework for how AI should be used and inventory every application that includes AI functionality.

For RIAs, this inventory should not be limited to obvious standalone AI tools. It should include:

- General AI platforms such as ChatGPT, Claude, Gemini, Copilot, and Grok
- AI-enabled CRM, portfolio management, communication, surveillance, cybersecurity, and document management tools
- AI note-taking, transcription, meeting summary, and task automation tools
- Vendor systems that have added embedded AI functionality
- Browser extensions and plug-ins
- Personal subscriptions used for firm business
- Internally developed automation or agentic workflows
- Tools used by contractors, consultants, outsourced CCOs, technology vendors, and independent representatives

The compliance challenge is that AI may already be present even when the firm has not formally adopted it. This is why defensible AI requires more than a policy statement. It requires discovery, classification, and control.

The Regulatory Question Is Not Whether AI Is Allowed

Regulators are not necessarily opposed to AI adoption. In fact, regulators are also using AI and expect the industry to explore it. The challenge is whether firms adopt AI in a way that protects investors and preserves accountability.

This distinction matters. A firm does not become defensible by refusing to innovate. It becomes defensible by showing that innovation is governed by a reasonable process.

AI should be introduced gradually, with clear use cases and defined controls. Firms should avoid implementing AI wholesale across the organization without understanding the operational, data, and supervisory implications. A measured rollout allows leadership to understand the impact of AI before it becomes embedded in critical workflows.

For small and mid sized RIAs, this is especially important. These firms often do not have large technology, compliance, and legal teams. A phased approach allows the firm to start with lower-risk internal use cases, such as summarizing public regulatory updates or organizing internal notes, before moving into higher-risk functions such as client communications, investment recommendations, policy drafting, compliance testing, surveillance, or data analysis involving client information.

The firm should define the purpose of each approved AI use case before deployment. It should also define how success will be measured. AI is not inherently valuable because it is new. It is valuable only if it improves a process without creating unmanaged risk.

Defensible AI Requires Data Awareness

The most important risk category for RIAs is data. Before approving an AI tool, a firm should understand where its data currently resides, how it is protected, and what happens to that data when it is entered into an AI system.

This is not just a cybersecurity question. It is a compliance question, a privacy question, and a fiduciary risk question. Client data, non-public personal information, account details, financial plans, investment holdings, internal communications, trading records, and supervisory materials may all create regulatory exposure if they are mishandled.

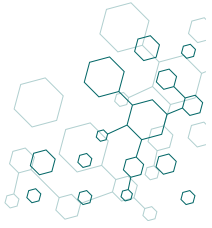

Firms must ensure that AI systems protect firm and client data from inadvertent disclosure. That requirement should become a core principle of AI governance. Before any AI tool is approved, the firm should understand:

- What data the tool can access
- Whether prompts and outputs are stored
- Whether data is used to train public or third-party models
- Where data is processed and stored
- Whether the vendor can access the data
- Whether data can be deleted or exported
- Whether the firm can retrieve usage logs
- Whether the tool supports enterprise controls
- Whether the tool integrates with existing cybersecurity and recordkeeping systems

A defensible AI program does not treat all AI tools equally. A public AI chatbot used with no client data presents a different risk profile than a vendor platform connected to email, CRM, portfolio data, or supervisory records. The firm should tier AI tools by risk and apply controls accordingly.

Core Principle

AI governance starts with inventory. A firm cannot supervise, test, restrict, or defend AI use it does not know exists.

SurgeONE.ai
AI-Powered. Practitioner-Led

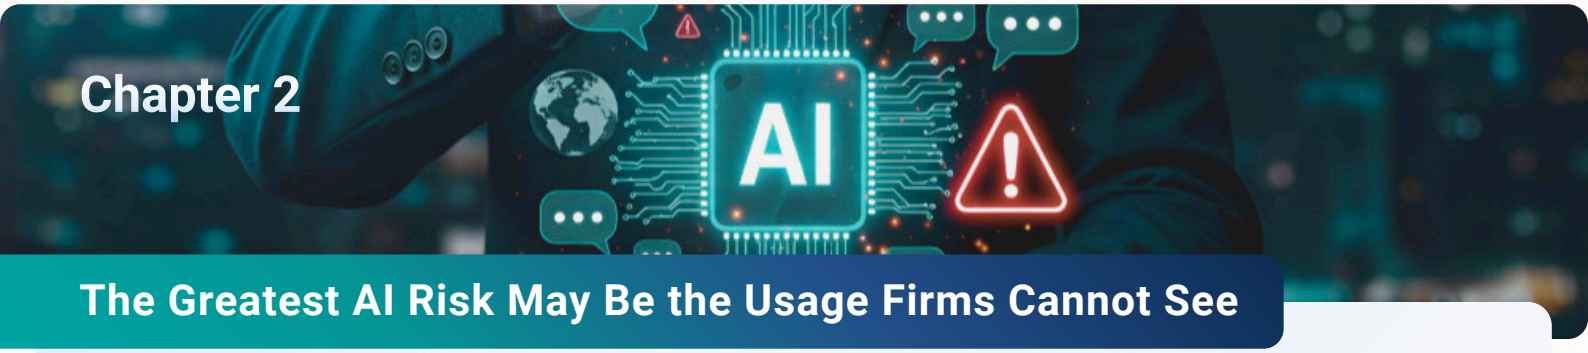

The Greatest AI Risk May Be the Usage Firms Cannot See

Many compliance risks arise from formal firm decisions. AI is different because a large portion of risk may arise from informal employee behavior. This is the problem of shadow AI.

Shadow AI occurs when employees, advisors, contractors, or representatives use unapproved AI tools for firm business outside the firm's governance framework. It may involve a personal ChatGPT account, a free AI assistant, a browser plug-in, a transcription tool, a personal Copilot account, or an AI feature inside a consumer application.

The exposure is especially significant for firms with decentralized workforces, remote employees, independent contractors, or advisors using their own devices and networks. Historically, firms have often relied on questionnaires asking independent representatives whether their devices were patched, encrypted, and free of unauthorized applications. The weakness of that model is obvious: firms have to rely heavily on self-reporting.

AI creates a similar challenge, but with broader implications. An employee can copy client data into a public AI tool in seconds. A representative can use a personal AI account to generate a client report. An advisor can paste account information into a chatbot to prepare a meeting summary. A contractor can use an AI note-taker without understanding whether the transcript is being stored or used to train a model.

If the firm does not know this activity is happening, it cannot evaluate privacy, recordkeeping, supervision, or data security implications.

The ability to detect shadow AI use is becoming one of the most important practical challenges for compliance and technology leaders.

Attestations Are Not Enough

Firms may be tempted to solve shadow AI with an annual certification. Employees can be asked to confirm that they are not using unauthorized AI tools or entering client data into public platforms. Certifications are useful, but they are not sufficient.

Regulators have already shown skepticism toward compliance programs that rely entirely on certifications without testing. The same concept applies to off-channel communications, outside storage, personal devices, and other areas where employee behavior can occur outside firm-approved systems. A certification may help show that expectations were communicated, but it does not prove that the firm had a reasonable supervisory process.

A defensible AI control framework should include multiple touchpoints:

- Policy prohibitions and permitted use standards
- Training on what is and is not allowed
- Employee certifications
- Vendor and application inventory
- Technical monitoring where feasible
- Reviews of metadata, documents, communications, and system logs
- Exception reporting
- Escalation and remediation procedures
- Periodic testing of whether controls are working

The firm should not rely on a single control. A policy without training will not be enough. Training without testing will not be enough. Testing without documentation will not be enough. Defensibility comes from the combination.

Employees Need Practical Training, Not Abstract Warnings

AI training should not be generic. Employees need to understand the specific types of information that should not be entered into unapproved tools. They need concrete examples of prohibited conduct.

For example, a firm should train employees not to paste the following into public or unapproved AI systems:

- Client names and contact information
- Account numbers
- Holdings or balances
- Financial plans
- Risk profiles
- Tax information
- Social Security numbers
- Internal supervisory notes
- Compliance reviews
- Trade blotters
- Due diligence files
- Confidential firm strategy
- Non-public business information

Employees may not understand what qualifies as non-public information or personally identifiable information unless the firm explains it clearly. Firms must train staff on what PII means, what confidential information means, and why client or firm information cannot be entered into unapproved AI tools.

This training should be role-specific. A portfolio manager, client service associate, operations analyst, CCO, CIO, and CEO may each use AI differently. The firm should explain AI risk in the context of real workflows.

For example:

- A client service associate should understand that using AI to draft a client email may create risk if the prompt includes client-specific financial information.
- A compliance officer should understand that using AI to draft policies may be helpful, but only if the output is reviewed against the firm's actual practices.
- A CIO should understand that an AI vendor's enterprise controls, data retention policies, and integration architecture are central to the firm's risk posture.
- A CEO should understand that productivity gains do not eliminate the need for supervision, accountability, and documentation.

SurgeONE.ai
AI-Powered. Practitioner-Led

Shadow AI Can Shift from Mistake to Intentional Circumvention

Training and documentation also matter because they help distinguish between an employee who made a mistake and an employee who intentionally circumvented firm controls.

If an individual goes around the firm's systems, sets up a separate account, or tries to avoid monitoring, the behavior can move from negligence to intentional avoidance.

This distinction is important. A firm cannot prevent every rogue action. But it can reduce firm-level exposure by showing that it had reasonable procedures, communicated them clearly, trained employees, collected certifications, tested compliance, and used reasonable methods to detect violations.

A defensible AI program should therefore document not only what the firm prohibits, but how the firm communicates, monitors, tests, and enforces those prohibitions.

CCO Question

If an examiner asked how your firm knows employees are not entering client information into unapproved AI tools, would your answer rely only on policy language and annual attestations?

SurgeONE.ai
AI-Powered. Practitioner-Led

Chapter 3

Human Oversight Is Not Optional. It Must Be Qualified, Documented, and Repeatable.

AI can accelerate work, but it cannot assume regulatory accountability. For RIAs, this point should be non-negotiable. Human oversight is not a symbolic control. It is the bridge between AI productivity and regulatory defensibility.

AI systems can hallucinate, produce inaccurate summaries, rely on outdated information, create biased outputs, misunderstand prompts, omit important context, or draft procedures that sound plausible but do not match the firm's actual business. These risks are amplified when AI is used in compliance, supervision, investment analysis, client communications, or policy drafting.

AI output should not be treated as final. It should be treated as work product requiring review.

The Human Must Be the Right Human

A common mistake is to define human oversight too broadly. Having a human in the loop is not enough if the reviewer lacks the expertise to evaluate the output.

For compliance-sensitive use cases, the reviewer must be qualified. If AI drafts a policy, the reviewer must understand the applicable regulatory requirements, the firm's business model, the firm's actual practices, and the consequences of adopting procedures that are inaccurate or unrealistic. If AI summarizes a regulatory issue, the reviewer must be able to identify omissions or misstatements. If AI supports surveillance, the reviewer must understand what the system is flagging and what it may be missing.

The human review process should answer four questions:

1. Is the AI output accurate?
2. Is it complete enough for the intended use?
3. Does it align with firm practices and regulatory obligations?
4. Has the review and approval been documented?

The fourth question is often the difference between a good internal process and a defensible one. If a firm cannot show who reviewed AI output, what they reviewed, what changes they made, and when they approved it, the firm may struggle to demonstrate effective supervision.

SurgeONE.ai
AI-Powered. Practitioner-Led

AI-Drafted Policies Create a Specific Risk

Using AI to draft compliance policies and procedures can be helpful, but it also creates a predictable risk. The output may sound sophisticated while failing to reflect the firm's actual operations.

Firms often have procedures that do not line up with their practices. Once those procedures are adopted, failure to follow them creates regulatory deficiencies.

This is an important point for CCOs. A policy is not defensible because it is well written. It is defensible because it accurately reflects what the firm does, what the rules require, who is responsible, how controls operate, and how exceptions are handled.

AI can generate generic procedures quickly. But generic procedures can be dangerous if they create obligations the firm does not actually meet. For example, if AI drafts a procedure stating that all AI outputs are reviewed weekly by compliance, but the firm does not perform that review, the firm has created a gap between written procedures and actual practice.

That gap can become an examination finding.

A defensible approach to AI-assisted policy drafting should include:

- Human review by a qualified compliance professional
- Alignment with the firm's actual business model
- Confirmation that assigned responsibilities are realistic
- Review of recordkeeping requirements
- Testing of whether procedures can actually be followed
- Version control and approval records
- Periodic updates as AI use cases evolve

AI can assist the drafting process. It should not own the compliance judgment.

AI Is Better Treated as an Analyst Than a Decision Maker

AI is best understood as a tool that can gather information, organize it, summarize it, and present it in useful formats. It can function like an analyst or paralegal, but it should not replace the professional judgment of the person responsible for the final decision.

This is a useful operating model for RIAs. AI can help summarize, classify, draft, organize, identify patterns, and accelerate analysis. But final judgment should remain with accountable professionals.

That principle is particularly important for:

- Compliance policies
- Regulatory responses
- Client-facing communications
- Investment recommendations
- Risk assessments
- Exception reviews
- Cybersecurity incident analysis
- Vendor due diligence
- Supervisory decisions
- Books and records determinations

The more a use case affects investors, client data, regulatory obligations, or firm supervision, the stronger the human oversight should be.

Accountability Cannot Be Outsourced to AI

Regulators will hold someone accountable when AI causes harm. The firm may be accountable as an entity. Individuals may also be accountable if they were responsible for oversight, approval, supervision, or implementation.

This principle should be reflected directly in AI governance documents. The firm should define who owns AI risk at the leadership level, who approves AI tools, who monitors ongoing use, who reviews outputs, who handles exceptions, and who escalates incidents.

For small and mid sized RIAs, the structure does not need to be overly complex. But it does need to be explicit.

At a minimum, firms should define:

- Executive sponsor for AI governance
- Compliance owner for AI policy and supervision
- Technology owner for security, access, logging, and vendor controls
- Business owner for each approved AI use case
- Reviewer or approver for high-risk outputs
- Incident owner for breaches, misuse, or control failures

The structure should match the firm's size, but the accountability should be clear.

Expert Quote

"If you use it, don't tell me later, 'Well, that wasn't on me. That was on AI."

SurgeONE.ai
AI-Powered. Practitioner-Led

Chapter 4

Regulators May Not Need New AI Rules to Enforce AI Failures

A common misconception is that firms can wait for specific AI rules before building governance. That is a risky assumption. The absence of AI-specific rules does not mean the absence of regulatory exposure.

A useful comparison is algorithmic and high-frequency trading. When algorithmic trading accelerated, regulators did not need an entirely new framework to bring enforcement actions. They applied existing rules related to manipulation, supervision, and firm responsibility.

The same logic applies to AI. If AI contributes to a misleading recommendation, a privacy breach, deficient supervision, inaccurate records, fraudulent representation, or failure to follow procedures, regulators can evaluate the conduct under existing obligations.

For RIAs, this means AI governance should be mapped to existing compliance categories rather than treated as a separate technology issue.

AI and Investor Harm

Regulatory scrutiny will focus heavily on investor harm or potential investor harm. This can arise in several ways.

AI may generate inaccurate client communications. It may summarize investment risks incorrectly. It may produce a recommendation or analysis based on outdated or incomplete data. It may expose client information. It may create biased outputs. It may be used in workflows that affect trading, supervision, or client service without adequate review.

The key issue is not whether the firm intended harm. The issue is whether the firm had reasonable controls to prevent, detect, and correct foreseeable risk.

If AI is used in any process that touches investor outcomes, firms should apply heightened scrutiny. That includes client communications, recommendations, portfolio reviews, risk assessments, financial plans, account surveillance, complaint analysis, trade monitoring, and marketing materials.

AI and Reg BI, Fiduciary Duty, and Recommendations

AI can create issues if its output constitutes or supports a recommendation to buy or sell particular securities.

For RIAs, the analysis should be tied to fiduciary duty. If an AI tool contributes to advice, recommendations, portfolio construction, or client-specific analysis, the firm must ensure the output is appropriate for the client's circumstances and consistent with the firm's obligations. AI should not be allowed to generate client-facing advice without qualified review.

Firms should be careful not to let productivity tools drift into advice functions. A tool originally approved for summarization may later be used to draft recommendations. A tool approved for research may later be used to produce client-ready commentary. A tool approved for internal analysis may become part of a client deliverable.

This is why use-case boundaries matter. AI governance should specify not only which tools are approved, but what those tools are approved to do.

AI and Privacy

Non-public personal information and client data exposure are among the most serious AI risks. Privacy risk is especially acute when employees use public AI tools or personal accounts. If client data is entered into an open AI system, the firm may lose control over where that data goes, how it is stored, whether it is retained, and whether it can be retrieved or deleted.

The firm's AI policy should include bright-line rules about data entry. These rules should be practical enough that employees understand them.

For example:

- Do not enter client names, account information, financial data, or personally identifiable information into unapproved AI tools.
- Do not upload client files, statements, reports, agreements, or compliance records into unapproved AI tools.
- Do not use personal AI accounts for firm business.
- Do not use AI note-taking tools for client or internal meetings unless approved.
- Do not use AI-generated output in client-facing communications without required review.
- Do not rely on AI for final compliance, legal, investment, or supervisory judgment.

The firm should then test whether employees are following those rules.

AI and Recordkeeping

AI also creates recordkeeping issues. If employees use AI to generate content, summarize meetings, draft communications, or analyze records, the firm must determine whether prompts, outputs, approvals, edits, and final versions need to be retained. If AI tools operate outside approved systems, the firm may lose records that should have been preserved.

This is another reason personal AI accounts are problematic. Even if no client data is exposed, the firm may not be able to retain or supervise the activity.

A defensible AI framework should define recordkeeping standards for each use case. For high-risk uses, the firm should preserve enough information to reconstruct what happened.

That may include:

- Prompt or instruction
- Data source used
- AI-generated output
- Human edits
- Reviewer name
- Approval date
- Final version
- Related client or account context, if applicable
- Exception notes
- Testing results

Regulatory Insight

The absence of AI-specific rules is not a safe harbor. Existing rules on supervision, privacy, records, fraud, fiduciary duty, and investor protection remain fully relevant.

The goal is not to retain unnecessary noise. The goal is to preserve evidence of supervision and decision-making where regulatory risk exists.

AI and Hallucinations

AI hallucinations are not just a technical issue. They are a compliance issue when AI output is used in regulated contexts.

A hallucination in a public blog draft may create reputational risk. A hallucination in a client report, compliance procedure, regulatory response, or supervisory review can create regulatory risk. The firm should assume AI output requires validation, especially when it involves legal, regulatory, investment, or client-specific content.

Testing should be ongoing. AI tools change. Vendor models update. Data sets become stale. Software changes may alter outputs. A control that worked six months ago may not work today.

SurgeONE.ai
AI-Powered. Practitioner-Led

Chapter 5

Vendor Oversight Is Now AI Oversight

Many RIAs will not build their own AI systems. They will adopt AI through vendors. That may include compliance platforms, CRM systems, portfolio tools, cybersecurity tools, surveillance systems, marketing platforms, document systems, research tools, meeting tools, and workflow automation systems.

This creates a major governance issue: the firm may rely on AI functionality embedded in vendor platforms without fully understanding how it works, what data it uses, or what risks it introduces.

A vendor may have been in place for many years under a contract that never contemplated AI. But that same vendor may now have AI tools embedded in its system.

This is a practical and immediate concern. AI functionality can be introduced through product updates, feature releases, integrations, or optional modules. The firm may not view this as a new vendor relationship, but the risk profile may have changed.

Contracts Need to Address AI-Specific Risk

Traditional vendor agreements may not adequately address AI. They may include confidentiality, data protection, service levels, and limitation of liability language, but not the specific issues created by AI models, prompts, outputs, training data, hallucinations, explainability, data retention, or human review.

Vendor contracts are becoming an active area of liability and negotiation. Vendors may try to tighten contract language to limit responsibility, while firms may demand more clarity about how the tool works and how risk is controlled.

RIAs should review vendor contracts with AI-specific questions in mind:

- Does the vendor use AI in the service?
- Is AI optional or embedded?
- What data can the AI access?
- Are prompts or outputs retained?
- Is firm or client data used to train models?
- Are public models, private models, or third-party models involved?
- Can the firm disable AI features?
- Does the vendor provide usage logs?
- Does the vendor support record retention?
- How does the vendor test for accuracy, bias, hallucination, and degradation?
- What human review does the vendor perform, if any?
- What liability does the vendor accept for AI-related failures?
- What notification obligations apply in the event of a breach or AI malfunction?
- What subcontractors or model providers are involved?
- What happens to data when the relationship ends?

These questions should be part of vendor due diligence and ongoing vendor oversight.

The Regulated Firm Remains Accountable

Regulators tend to pursue the regulated entity when vendor systems fail in ways that affect regulated obligations. The same principle should guide AI oversight. A firm cannot assume that vendor use of AI shifts responsibility away from the firm. If the vendor's AI tool causes a supervisory failure, inaccurate output, data breach, or recordkeeping deficiency, the firm may still be asked why it approved the vendor, how it reviewed the tool, what controls it required, and how it monitored performance.

Vendor due diligence should therefore be risk-based. High-risk AI vendors should receive deeper review. A vendor that uses AI to summarize public content is not the same as a vendor that has access to client data, communications, compliance records, trading information, or supervisory workflows.

A risk-tiering model might consider:

- Type of data accessed
- Client impact
- Regulatory impact
- Degree of automation
- Whether output is client-facing
- Whether output affects advice or supervision
- Whether the vendor is critical to operations
- Whether the tool is explainable and testable
- Whether the firm can monitor usage
- Whether the firm can preserve records
- Contractual protections
- Vendor financial and operational resilience

SurgeONE.ai
AI-Powered. Practitioner-Led

Vendor AI Can Create Economic and Liability Tension

AI vendors may promise to reduce headcount, streamline operations, and allow a smaller team to do the work of a larger one. But if vendors are asked to accept more liability for AI-driven errors, pricing and contract terms may change.

This matters for firm leadership. AI should not be evaluated only as a cost-saving tool. It should be evaluated as a risk-transfer and control issue. If a vendor promises major efficiency gains, the firm should ask what risk remains with the firm and what controls are necessary to make the workflow defensible.

In practice, firms should avoid replacing human judgment in high-risk functions without carefully assessing whether the AI system can be supervised, tested, explained, and documented.

Ongoing Due Diligence Is Required

AI vendor review should not be a one-time procurement exercise. AI systems change frequently. Software updates, model changes, new features, new data flows, and new integrations can alter risk.

AI is not a plug-and-play investment. For RIAs, ongoing due diligence should include periodic vendor reviews, contract updates, control testing, incident tracking, feature review, and confirmation that approved use remains consistent with the firm's policy.

The firm should also require vendors to notify it of material AI-related changes. These may include:

- New AI features
- Changes in model providers
- Changes in data usage
- Changes in retention policies
- New integrations
- Material incidents
- Changes in subcontractors
- Changes in security controls
- Changes in output review processes

Vendor Governance Question

Has your firm reviewed whether long-standing vendors have added AI capabilities that were not addressed in the original contract?

The goal is to avoid a situation where a vendor relationship originally approved for one risk profile quietly evolves into something materially different.

Conclusion

The Practical Path Forward: Use AI, But Make It Defensible

AI will become more embedded in financial advisory firms. Avoiding AI entirely may not be practical or competitive. But adopting AI without governance creates exposure that can accumulate quickly and quietly.

The responsible path is not fear. It is structure.

Firms should use AI, but not blindly. They should keep experts in the loop, train their workforce, and implement active controls to test, monitor, and certify results. Examiners are likely to ask not only whether AI is being used, but how the firm governs, supervises, and documents that use.

That operating model should guide every RIA.

A defensible AI program should include seven foundational elements.

1. Inventory - The firm should identify all AI tools currently in use, including standalone platforms, vendor-embedded AI, personal accounts, browser extensions, note-taking tools, and internal automations. Inventory should include who uses each tool, what data it accesses, what business purpose it serves, and whether it is approved.

2. Policy - The firm should define permitted, restricted, and prohibited AI uses. The policy should address data entry, client information, personal accounts, vendor tools, recordkeeping, human review, escalation, and sanctions for circumvention.

3. Training - Employees need practical training on what AI tools can and cannot be used for, what information cannot be entered into unapproved systems, what human review is required, and why violations create risk for clients, the firm, and the individual.

4. Qualified Human Review - The firm should require qualified review for AI output used in compliance, supervision, client communications, investment analysis, policies, regulatory responses, or other high-risk workflows. The review should be documented.

5. Vendor Oversight - AI vendor due diligence should be risk-based and ongoing. Contracts should address data use, model training, retention, logs, security, liability, incident notification, and the firm's ability to supervise and preserve records.

6. Testing and Monitoring - The firm should not rely solely on attestations. It should test whether AI policies are being followed, monitor for shadow AI where feasible, review metadata and system activity where appropriate, and periodically evaluate whether approved tools remain accurate, current, and fit for purpose.

SurgeONE.ai
AI-Powered. Practitioner-Led

7. Documentation - The firm should document its AI governance decisions, approvals, training, certifications, testing, exceptions, vendor reviews, incidents, and remediation. Documentation is what allows the firm to show reasonable effort, reasonable supervision, and reasonable due diligence.

AI may save time. But without governance, today's efficiency can become tomorrow's examination issue. The firms that benefit most from AI will not be the ones that adopt it casually. They will be the ones that adopt it deliberately, align it with regulatory obligations, preserve human accountability, and build controls that can withstand scrutiny.

Defensible AI is not a technology purchase. It is a compliance operating discipline.
And for RIAs, that discipline should begin now.

A Practical AI Governance Checklist for RIAs

☐	Inventory - Has the firm identified all AI tools in use, including vendor-embedded AI and personal AI accounts used for business?
☐	Data Protection - Has the firm clearly defined what information may never be entered into public or unapproved AI tools?
☐	Policy - Does the firm's AI policy distinguish between approved, restricted, and prohibited use cases?
☐	Human Review - Are qualified individuals reviewing high-risk AI outputs before use?
☐	Recordkeeping - Does the firm know which prompts, outputs, approvals, and final materials must be retained?
☐	Vendor Oversight - Have vendor contracts been reviewed for AI functionality, data use, liability, retention, and incident notification?
☐	Testing - Does the firm test whether employees are following AI policies, rather than relying only on certifications?
☐	Incident Response - Does the firm know what it would do if client data were entered into an unapproved AI tool?
☐	Documentation - Could the firm show an examiner evidence of AI governance, training, approvals, testing, and remediation?

www.surgeone.ai info@surgeone.ai +1 650-294-8889