A CCO's Guide to AI

Tools, Tips, and Compliance Best Practices

AI is rapidly changing and RIAs are quickly beginning to utilize it for a wide range of business activities. This puts CCOs in a tough position: balancing the potential benefits of AI for their firms with the need to ensure its compliant use.

This guide is intended to help CCOs understand the most common types of AI tools their teams are likely using, the key risks associated with each, and practical recommendations for building a compliant AI framework.

Table of Contents:

I. Types of AI Tools CCOs Should Know

Not all AI tools carry the same risks. Understanding the distinctions between tool types is the first step toward building appropriate guardrails.

AI Connectors

AI connectors are third-party tools that link Generative AI platforms (like ChatGPT or Claude) to your external services — email, cloud storage, CRMs, and more. It's important to note that adviser regulatory obligations do not transfer to the AI vendor.

For example, Smarsh and Global Relay have added AI connectors to monitor communications and flag regulatory exceptions. SharePoint can be connected to an AI tool that searches, summarizes, and drafts reports from firm documents. These integrations can be powerful, but they require the same risk assessment and oversight as any new capability.

> **Example:** An adviser connects Salesforce to Claude. After a client meeting, the AI drafts a follow-up email by reading meeting notes logged in the CRM. Key questions: Was the summary reviewed before sending? Was it captured as business communication for recordkeeping?

AI Connector and App Risks

1. **Unapproved Data Flows:** Connecting AI to CRM, email, or document systems can expose sensitive data to third-party vendors without adequate review.
2. **Implicit Approval Assumptions:** Employees may be using AI features embedded in approved platforms without realizing that those features have not been independently vetted.
3. **Training Data Exposure:** Many AI tools — especially free or lower-tier versions — use firm and client data to train their models. This data may then surface in other users' outputs.

CCOs should apply standard controls: risk assessments, written policies and procedures, employee training, and ongoing supervision.

Apps Adding AI Features

Many existing platforms like CRMs, document management systems, and communication tools are quietly adding AI features. Employees may assume that, because a platform is already approved, its new AI features are approved too. This is a common and risky assumption.

Before using any new AI feature within an existing platform, firms should:

- Carefully review updated contracts, terms of service, and data processing agreements.
- Confirm whether firm or client data may be used to train AI models.
- Determine whether the AI feature can be disabled if it has not been reviewed.
- Apply the same approval and due diligence process as standalone AI tools.
- Be aware that embedded or "free" AI features often use your data to train their models and may not advertise this fact.

Generative AI vs. Agentic AI

Understanding the distinction between these two types of AI is essential for calibrating your firm's risk posture.

Generative AI

How it works: A human types a prompt, the AI produces output (a draft email, a summary, a marketing piece), and a human reviews and approves it before any action is taken. This is a relatively low-risk use of AI because the human remains the decision maker throughout.

Agentic AI

How it works: Agentic AI has more autonomy. Agentic AI is triggered by an event (a file dropped into a folder, an email received) and makes sequential decisions before a human ever sees a result. This part is key: Agentic AI can trigger a series of actions without review, consent, or even knowledge of the actions.

With Agentic AI, it is absolutely critical that firms not only understand how it works, but that they also understand its capabilities, so that they have proper guardrails in place.

CCO Considerations for Agentic AI:

- Require documented approval before any agentic workflow goes live.
- Establish monitoring to verify the agent is performing as intended.
- Conduct periodic audits to catch drift or unintended behavior.
- Note that tools like Claude Cowork are significantly more powerful than browser-based versions — they can access files, shared drives, browser sessions, and connected applications.

> **Key Principle:** AI can make things go very well — or very badly — very fast. Agentic AI in particular warrants more restrictive approval processes and robust monitoring.

AI Notetakers

AI notetakers are among the most common AI tools used by RIAs, and they raise a distinct set of regulatory questions. These tools automatically join meetings, transcribe conversations, identify speakers, extract action items, and store recordings — all in real time.

Popular examples include Otter.ai, Fireflies.ai, Fathom, and Notion AI. They typically integrate directly with Zoom, Google Meet, or Teams and appear as a bot participant on the call.

It's important to note that even AI platforms that claim they do not "record" calls must still listen to produce a transcription.

AI Notetaker Risks

1. **Consent and Transparency:** Most states operate under a one-party consent, but approximately 25% of states require two-party (or all-party) consent before a conversation may be recorded or transcribed. Advisers must understand the consent laws in the states where their clients are located and be fully transparent when using these tools. Verify the current two-party consent states before implementing any firm-wide policy, as state laws continue to evolve.
2. **Books and Records:** If a transcript is used to document investment instructions or decisions, it is a record subject to retention requirements. If a summary is sent to meeting attendees, it is written business communication that must be retained under the Books and Records Rule. Firms must decide whether it is operationally simpler to retain everything captured by AI tools or to retain selectively — and accept the risk of gaps.
3. **Data Storage and Reg S-P:** AI tools may capture personal information, MNPI, and sensitive client data. Under Reg S-P, this information must be protected. The location of stored records must be evaluated carefully. AI tools that train on your data may expose confidential information in public systems.
4. **Discovery Risk:** Any communication that is recorded or written down can be discovered by regulators, in litigation, or in examinations. CCOs should ask: Are we comfortable with a regulator reviewing a transcript of every call we conduct?

II. Frequently Asked Questions

1. **Must AI transcripts and recordings be retained?**
   RIAs do not have the same recording obligations as broker-dealers under FINRA's taping rule or Rule 17a of the Exchange Act. The operative question is whether the AI-generated content (**such as** a transcript, summary, or recording) is being sent to a client as a written business communication or **being** used to satisfy any other recordkeeping requirements. If yes, retain it. Separately, firms should weigh whether archiving everything is simpler and less risky than selective retention.
2. **How do I obtain consent for AI notetaking?**
   In practice, a common approach is a verbal disclosure at the start of a call (for example, "This meeting is being recorded or transcribed using an AI tool") or a written notice included in the meeting invitation, **or both**. If a participant does not consent, the meeting should not be transcribed and the adviser must take notes through other means. **Consent requirements vary depending on state law. Questions regarding specific state laws and methods of obtaining consent should be directed to legal counsel.**
3. **Are AI-generated summaries covered by the Books and Records Rule?**
   They can be. Written business communication is required to be retained. If an AI tool generates a summary that is shared with a client, or if the firm distributes a transcript, that content is arguably written communication subject to retention obligations.
4. **What if the AI transcript contains errors?**
   Whoever sends or causes the communication to be sent is responsible for its accuracy. AI hallucinations (**such as** errors in investment timing, amounts, risk ratings, or transaction costs) can cause material harm. A human review step is essential before any AI-generated content is shared externally.
5. **How do I supervise AI use across the firm?**
   Start by establishing an approved list of AI tools and an associated written policy. The next operational risk is unapproved AI usage — an employee installing a consumer-grade tool on a firm device, or using a free version that transmits sensitive data to public systems. Assume that many employees are already comfortable using AI in their daily work and build your framework accordingly.

III. Recommendations for Implementing AI

Vendor Due Diligence

- Conduct thorough vendor due diligence on all AI products before approving use.
- Include AI usage questions in all vendor due diligence reviews, even for non-AI vendors whose platforms may be adding AI features.
- Review updated contracts, terms of service, and data processing agreements whenever a vendor launches a new AI feature.
- Confirm whether the AI vendor may use firm or client data to train its AI models.
- Look for applicable protections: data processing agreements, no-training clauses, SOC 2 compliance.

Platform and Subscription Management

- Use the highest enterprise subscription tier available for AI platforms as these typically offer stronger data protections and contractual commitments.
- Disable AI add-ins and embedded features in existing platforms unless they have been formally reviewed and approved.
- Maintain an approved list of AI tools and a corresponding list of prohibited tools and platforms.

Employee Training and Supervision

- Conduct thorough and routine employee training on approved AI tools, proper usage, and associated risks.
- Establish a clear process for employees to request approval of new AI tools before use.
- Build monitoring into your supervisory framework to detect unapproved AI usage.

Rollout and Governance

- When introducing new AI products, use a phased approach. Start with a small beta group, document learnings, and build best practices before broader rollout.
- Establish a process to monitor ongoing AI usage and audit agentic workflows periodically.
- Remain vigilant. AI capabilities and their associated risks are evolving rapidly.

SEC Exam Readiness: Regulators are increasingly focused on AI governance in examinations. Be prepared to demonstrate that your firm has written AI policies, a vendor due diligence process for AI tools, employee training records, and a supervisory framework that addresses both approved and unapproved AI usage.

Questions? We're Here to Help.

Fairview has a team of dedicated compliance and cybersecurity experts who are well-versed in AI usage, SEC expectations, and cybersecurity best practices. Whether you're evaluating which products are right for your firm or need guidance building a compliant AI framework from the ground up, we can help. Contact us today to set up a conversation.

About Fairview

Founded in 2005, Fairview®, LLC provides a full range of back-office support services for investment advisers and other financial institutions.
For more information, visit FairviewInvest.com.

FAIRVIEW