GUIDE

Top 10 AI Best Practices for Financial Services Compliance

smarsh

Artificial intelligence (AI) is reshaping financial services at a pace not seen since the introduction of the internet. FINRA has identified more than 4,000 AI use cases across client engagement, compliance, and trading.

As clients expect faster and more responsive services, firms are under pressure to adopt AI tools not only to showcase innovation but also to improve productivity and reduce costs. However, whether content is human- or AI-generated, compliance obligations remain the same and fall under existing regulatory rules, including those from the SEC, FINRA, and FCA.

The SEC 2026 Examination Priorities noted:

"The Division remains focused on registrants' use of certain products and services, such as automated investment tools, AI technologies, and trading algorithms or platforms, and the risks associated with the use of emerging technologies and alternative sources of data... The Division will assess whether firms have implemented adequate policies and procedures to monitor and/or supervise their use of AI technologies, including for tasks related to fraud prevention and detection, back-office operations, anti-money laundering (AML), and trading functions, as applicable."

The EU AI Act adds another challenge, requiring risk assessments, transparency, and human oversight for high-risk AI systems. Financial services firms must navigate this complex regulatory landscape while maintaining their innovation agenda.

The central compliance challenge is determining when AI outputs cross into regulated territory. While not every AI interaction creates a record, certain use cases — such as client-facing communications, pre-trade analysis, or KYC summaries — may trigger capture and supervision requirements.

This guide provides ten best practices to help firms adopt AI responsibly and balance innovation with compliance.

1 Establish a risk-based governance framework

Governance is the anchor that prevents AI adoption from drifting into risky territory. Without a structured framework that considers regulatory, information security, data privacy, and intellectual property risks, firms risk uneven adoption, inconsistent interpretations, and exposure to both regulatory and business threats.

In its 2026 Examination Priorities, the SEC emphasized that firms must address risks arising from “automated investment tools, AI technologies, and trading algorithms,” particularly where these tools influence client recommendations or investment strategies. Regulators expect governance frameworks to anticipate and manage these risks across business units.

A governance framework should include:

- **Use case risk assessments**
  High-risk scenarios include AI outputs that influence client decisions, investment strategies, or regulatory reporting. Lower-risk scenarios may include internal summarization or productivity enhancements.
- **A governance council**
  The council should involve compliance, risk, IT, legal, and business leaders — stakeholders that ensure balanced decision-making.
- **Risks mapped to existing rules**
  Identify what compliance rules already apply. For example, internal AI use for research may not be a record, but if outputs inform trading or advice, recordkeeping obligations may apply.

IMPORTANT NOTE:

Governance councils should include clear executive-level ownership and be operationalized into an ongoing program that can revisit project status, reassess risks, and respond to evolving user demands. When implementing AI, consider piloting in lower-risk scenarios before expanding to the rest of the firm.

2 Define and enforce AI usage policies

Without clear policies, employees will use AI inconsistently, often in ways that expose the firm to compliance and reputational risks. Policies establish accountability and create a baseline for training and enforcement.

Clear policies establish boundaries for employees and reduce the risk of shadow AI use.

Policies should consider:

- **Defining** approved AI tools and their intended use cases
- **Requiring** employee attestation of compliance
- **Prohibiting or monitoring** use of unapproved AI platforms and embedded features for business purposes

IMPORTANT NOTE:

Policies should emphasize that not all AI outputs are records. It depends on the context. Supervisors and compliance officers must evaluate use cases to determine when capture is required.

3 Build robust training and awareness programs

Employees often drive both the benefits and the risks of AI. Without training, they may misuse tools, misinterpret outputs, or expose the firm to unnecessary liability.

Training is essential for embedding compliance into AI adoption.

This may include:

- **Educating** employees on when AI outputs may constitute records
- **Highlighting** scenarios where AI misuse could create risk, even if outputs are not retained
- **Running** exercises simulating errors, such as an AI drafting a misleading client memo

IMPORTANT NOTE:

Training reinforces that employees cannot rely on AI blindly. Responsibility for compliance will always **remain with humans**.

4 Evaluate whether AI outputs constitute business records

Determining whether AI outputs are records is the heart of compliance. A misstep can mean either over-preserving irrelevant data (increasing costs and risks) or under-preserving regulatory records (leading to fines and enforcement).

Determining whether AI outputs are records is the heart of compliance. A misstep can mean either over-preserving irrelevant data (increasing costs and risks) or under-preserving regulatory records (leading to fines and enforcement).

A FINRA notice notes that depending on how they’re using generative AI, member firms could implicate virtually every area of a firm’s regulatory obligations.

The SEC's 2026 Examination Priorities highlight that the Division will "examine firms that engage in activities such as automated investment advisory services, recommendations, and related tools and methods," with particular focus on the risks associated with emerging technologies.

These areas often produce content or recommendations that may — depending on context and use — meet recordkeeping criteria under existing rules. Firms should evaluate and document their determinations accordingly.

The key compliance question remains: When does AI-generated output become a record? Regulators have not provided definitive answers, but firms have existing rules to inform their judgment.

Firms should consider the following to determine when an AI-generated output is a record:

- Was the output used in client communication?
- Did it inform an investment decision?
- Would it be retained if created by a human?

Examples of potential records:

- An AI-generated investment summary relied upon by a portfolio manager
- An AI-generated meeting note capturing client commitments

Examples that aren't records:

- A brainstorming draft used internally without informing decisions
- AI-generated first draft of an email that was substantially rewritten before delivery

IMPORTANT NOTE:

The determination ultimately rests with the firm's supervisory procedures, risk appetite, and engagement with regulators. Ultimately, if the content was created by an employee and it needed to be retained, the same applies to the AI-generated version of that same content. Firms should document their rationale for treating certain outputs as records — or not — so that positions are defensible in examinations.

5 Capture and supervise embedded AI features

As AI becomes embedded by default into enterprise platforms, firms risk missing records simply because employees are unaware that features are AI-driven. This makes supervision more complex.

Firms should:

- **Review** outputs from embedded features (auto-summaries, suggested replies) to determine whether they qualify as records
- **Capture** outputs that intersect with regulated activity, while documenting why others are not retained
- **Update** supervisory procedures to account for embedded AI

> **IMPORTANT NOTE:**
>
> Assessing embedded AI features will ensure consistency with regulatory expectations that rules apply regardless of how the content was generated.

6 Enhance data protection and privacy controls

Privacy breaches involving AI are costly and damaging. Regulators treat AI tools as extensions of a firm's data environment, meaning privacy standards apply equally.

AI can inadvertently expose sensitive data.

Firms should consider:

- **Preventing** client PII or MNPI from being entered into unsecured AI systems
- **Assessing** risks of AI features embedded in enterprise platforms like Bloomberg, Refinitiv, Microsoft Teams, or Slack
- **Deploying** data loss prevention and endpoint monitoring to detect unauthorized use

IMPORTANT NOTE:

Even if firms determine certain outputs are not records, they must still safeguard client information under privacy rules such as Reg S-P and GDPR.

7 Strengthen vendor risk management

Vendors often act as unseen participants in compliance workflows. If vendor practices are not carefully reviewed, firms risk exposing sensitive client data or relying on tools that fail regulatory standards.

FINRA states rules apply regardless of whether firms are directly developing generative AI tools for proprietary use or using third-party technologies.

AI tools provided by third-party vendors present unique risks. Firms must:

- **Conduct** due diligence on how vendors handle data, prompts, and outputs
- **Assess** whether vendor practices align with privacy regulations and firm standards
- **Require** contractual assurances against model training on firm data without consent

IMPORTANT NOTE:

Vendor assessments should also consider whether outputs generated through these tools could be deemed records, and if so, whether capture is feasible.

8 Embed human-in-the-loop oversight

AI tools can accelerate workflows, but they may lack context. Human oversight ensures that outputs are trustworthy, compliant, and aligned with client expectations.

AI cannot replace human judgment.

Firms should consider:

- **Requiring** supervisors to review outputs before they are used in regulated workflows
- **Training** employees to validate AI-generated content, ensuring it aligns with firm policies and regulatory standards
- **Emphasizing** prompt discipline, such as avoiding sensitive data in inputs and ensuring outputs are factually accurate

IMPORTANT NOTE:

Human oversight provides the safeguard regulators expect, especially when firms take the position that certain AI outputs are not records.

9 Integrate AI into supervision and surveillance programs

Supervision is not just about capturing records. It's about demonstrating reasonable oversight. Regulators expect firms to show visibility into how AI is used, even if not all outputs are retained.

The SEC has stated that it will "assess whether firms have implemented adequate policies and procedures to monitor and/or supervise their use of AI technologies, including for tasks related to fraud prevention and detection, back-office operations, AML, and trading functions." It will also review "firm integration of regulatory technology to automate internal processes and optimize efficiencies."

This reinforces the need for firms to demonstrate not only that AI is supervised, but that supervisory systems evolve alongside emerging technologies.

Even when AI outputs are not considered records, supervision may still be warranted.

Firms should:

- **Extend** existing supervisory review processes to AI-assisted content
- **Capture** and log AI use when it touches business communications
- **Implement** monitoring to detect shadow AI and unusual phrasing that may suggest AI involvement

IMPORTANT NOTE:

Regulators have stressed that supervision must be "reasonable and visible." The absence of capture should be defensible, supported by other safeguards such as training, monitoring, or technical controls.

10 Test, audit, and continuously improve controls

AI is non-deterministic and constantly evolving. Controls that are effective today may be obsolete tomorrow. Continuous auditing helps firms stay ahead of both risks and regulatory scrutiny.

AI governance must evolve with the technology.

Firms should consider:

- **Auditing** AI outputs to confirm compliance with stated policies
- **Conducting** shadow AI detection exercises to identify unapproved use
- **Reassessing** recordkeeping positions regularly, documenting rationales

IMPORTANT NOTE:

Testing and auditing not only strengthen controls but also provide evidence to regulators that decisions about record status are thoughtful and risk-based.

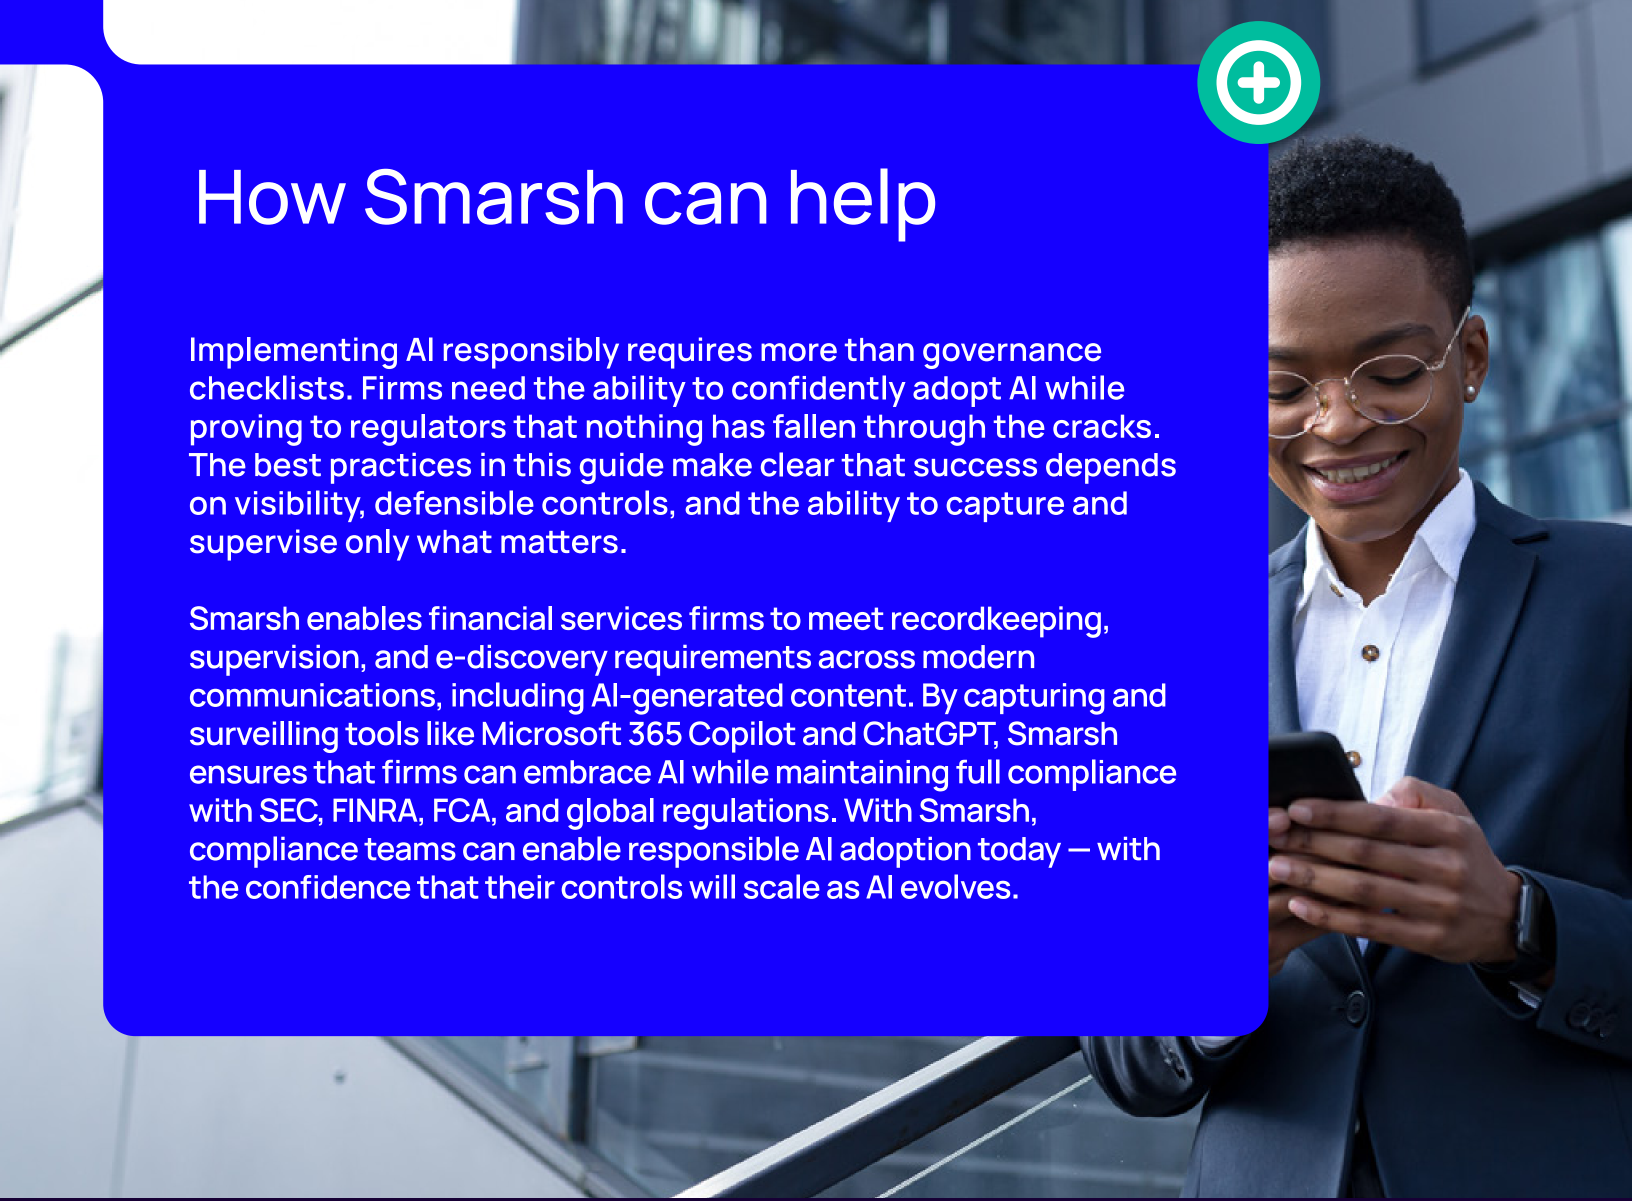

How Smarsh can help

Implementing AI responsibly requires more than governance checklists. Firms need the ability to confidently adopt AI while proving to regulators that nothing has fallen through the cracks. The best practices in this guide make clear that success depends on visibility, defensible controls, and the ability to capture and supervise only what matters.

Smarsh enables financial services firms to meet recordkeeping, supervision, and e-discovery requirements across modern communications, including AI-generated content. By capturing and surveilling tools like Microsoft 365 Copilot and ChatGPT, Smarsh ensures that firms can embrace AI while maintaining full compliance with SEC, FINRA, FCA, and global regulations. With Smarsh, compliance teams can enable responsible AI adoption today — with the confidence that their controls will scale as AI evolves.

smarsh

Smarsh® enables companies to transform oversight into foresight by surfacing business-critical signals from the most in-demand communications channels. Regulated organizations of all sizes rely upon the Smarsh portfolio of cloud-native digital communications capture, retention and oversight solutions to help them identify regulatory and reputational risks within their communications data before those risks become fines or headlines.

Smarsh serves a global client base spanning the top banks in North America, Europe and Asia, along with leading brokerage firms, insurers, and registered investment advisors and U.S. state and local government agencies. To discover more about the future of communications capture, archiving and oversight, visit www.smarsh.com

Smarsh provides marketing materials for informational purposes only. Smarsh does not provide legal advice or opinions.
You must consult your attorney regarding your compliance with applicable laws and regulations.

Guide - 12/25

☎ 1-866-762-7741 www.smarsh.com @SmarshInc SmarshInc Company/smarsh