PUBLICATION

FINRA's GenAI Playbook: Real Accountability for Broker-Dealers

Authors: Matthew George White, Lori H. Patterson

January 05, 2026

The Financial Industry Regulatory Authority's (FINRA's) 2026 Annual Regulatory Oversight Report (the Report) marks a notable escalation in the regulator's attention to generative artificial intelligence (GAI). While FINRA has been discussing AI for several years, its latest guidance reflects a clear pivot: GAI is no longer theoretical, experimental, or limited to innovation labs; it is increasingly embedded in day-to-day firm operations, and FINRA expects governance to keep pace.

The Report makes one thing unmistakably clear: If your firm is using GAI to draft communications, flag suspicious activity, draft policies and procedures, or streamline compliance workflows, regulators expect you to govern it with the same diligence you'd apply to your registered representatives, because in FINRA's eyes, the algorithm is now part of your supervisory chain and will be examined as such.

This article explores FINRA's heightened expectations for GAI governance, examines the specific risks that have caught regulators' attention, and provides a practical roadmap for broker-dealers navigating this new frontier of algorithmic accountability.

1. FINRA's New Focus: GAI Is "Just Technology" Until It Isn't

FINRA continues to emphasize that its rules are "technology neutral." In other words, existing obligations apply regardless of whether a task is performed by a human, traditional software, or an AI tool. But the Report clarifies that GAI introduces risk characteristics that demand heightened scrutiny, particularly where firms rely on AI outputs for regulated functions.

FINRA explicitly calls out potential implications under rules governing:

- Supervision (FINRA Rule 3110)
- Communications with the public
- Books and records
- Fair dealing and investor protection

If a firm is using GAI as part of its supervisory system, FINRA expects policies and procedures to address the integrity, reliability, and accuracy of the model itself, not just the end result. Put differently: "The AI did it" will not be a defense.

2. A Brief Look Back: FINRA's AI Guidance to Date

Understanding how FINRA arrived at this position requires a brief look at the regulator's evolving AI guidance. In fact, FINRA's latest commentary builds on several years of incremental guidance, including:

- Regulatory Notice 24-09, reminding firms that GAI and large language models do not alter regulatory obligations;
- Prior FAQs on advertising regulation, particularly where AI-generated content is used;

- FINRA's 2020 Report on Artificial Intelligence in the Securities Industry (which now reads as an early warning shot); and
- Joint investor alerts with the SEC and North American Securities Administrators Association (NASAA) addressing AI-enabled fraud.

What's different now is specificity and the intent behind it. Rather than speaking in abstractions, FINRA is cataloging real-world use cases, observed practices, and emerging patterns with examination-ready precision. The regulator has moved from learning mode to accountability mode. Translation: Firms should read this Report not as guidance, but as a preview of the questions examiners will be asking and the documentation they'll expect to see.

3. **What FINRA Is Actually Seeing Firms Do with GAI**

FINRA reports that GAI adoption among member firms is accelerating, particularly for internal efficiency and information management. In a rare glimpse into industry practices, the Report provides a de facto benchmarking guide showing firms not only what FINRA is watching, but what their competitors are already doing. The most common use case – by a wide margin – is summarization and information extraction, such as condensing lengthy documents or pulling key facts from unstructured data.

Other commonly observed use cases include:

- Conversational AI and internal chatbots
- Drafting reports, emails, and marketing materials
- Translation and transcription
- Coding assistance
- Workflow automation and process intelligence
- Pattern recognition and threat detection
- Synthetic data generation
- Personalization and recommendations

Notably, FINRA is also closely watching the emergence of AI agents – systems capable of autonomously planning and executing tasks across multiple systems with minimal human intervention.

The diversity of use cases presents both opportunity and risk. While firms are clearly finding value in GAI's efficiency gains, each application introduces distinct regulatory considerations, and FINRA expects governance frameworks to account for all of them. This leads to a potentially uncomfortable reality: If you think your firm isn't using GAI, you're probably wrong. And if you think FINRA hasn't noticed what you're using it for, you're definitely wrong.

4. **The Risks That Have FINRA's Attention (and Should Have Yours)**

FINRA does not mince words about the risk profile of GAI, repeatedly emphasizing issues that will feel familiar to cybersecurity and privacy professionals, but are now firmly on regulators' radar.

Core GAI Risks

- Hallucinations: Confidently delivered but inaccurate outputs that can distort compliance decisions, client communications, or supervisory reviews.
- Bias and data drift: Outputs influenced by skewed, outdated, or incomplete training data.

- Privacy and confidentiality: Risks of exposing sensitive or proprietary data through prompts, outputs, or vendor integrations.
- Cybersecurity exposure: Expanded attack surfaces, including prompt manipulation and AI-enabled threat activity. Practically, this means firms face dual exposure: unauthorized employee use of consumer AI tools (with, potentially, sensitive data leaving the firm's control) and increasingly sophisticated AI-powered attacks (such as GAI powered phishing content) for which traditional security controls and training may be insufficient.

Additional Risks from AI Agents

FINRA flags several agent-specific concerns, including:

- Excessive autonomy without human validation
- Agents exceeding their intended authority or scope
- Limited auditability and explainability
- Improper handling of sensitive data
- Misaligned incentives or reinforcement logic

In short, the more "independent" the AI, the greater the expectation that firms implement guardrails, logging, and human-in-the-loop oversight.

Bottom line: If an AI agent can act without human approval, your firm must be able to explain what decisions the agent made, why it made them, and how supervisory personnel would detect and correct errors quickly.

5. **FINRA's Practical Message: Govern First, Deploy Second**

FINRA stops short of prescribing a single AI governance model, but its expectations are unmistakable. Firms using – or considering using – GAI should be prepared to demonstrate all of the following.

Enterprise-Level Governance:

- Formal review and approval processes involving appropriate business and technology stakeholders
- Clear policies covering development, deployment, use, and monitoring of GAI
- Comprehensive documentation throughout the use case lifecycle

Risk Management and Testing:

- Pre-deployment (or verification thereof) testing for accuracy, reliability, integrity, and privacy
- Evaluation of third-party AI vendors as part of the firm's broader cybersecurity program, including proactive inquiry into whether vendors are using AI in their services, as many vendors have embedded GAI capabilities without explicit disclosure to clients

Supervision and Monitoring:

- Ongoing monitoring of prompts, outputs, and performance
- Logging and retention of prompts and responses for accountability
- Model version tracking and change management
- Human review of AI outputs, particularly in regulated workflows

FINRA's guidance aligns closely with emerging best practices under frameworks such as National Institute of Standards and Technology's (NIST's) AI Risk Management Framework, but with a distinctly securities-law overlay.

Key Takeaways for Firms

FINRA's message is ultimately pragmatic: Innovation is permitted, but abdication of responsibility is not. Firms are free to deploy GAI tools but remain fully accountable for the outcomes.

For broker-dealers and other regulated financial institutions, now is the time to:

- Inventory current and planned GAI use cases;
- Map those uses to supervisory, recordkeeping, and communications obligations;
- Update AI governance and vendor risk frameworks; and
- Prepare for examiners who will increasingly ask, "Who's supervising the algorithm?"

Because when AI makes a mistake, FINRA will still expect a human answer.

If you have questions about FINRA's GAI guidance, AI governance frameworks, or how these expectations intersect with cybersecurity, data privacy, or regulatory compliance, please contact Lori Patterson, Matt White, or any member of Baker Donelson's Broker-Dealer/Investment Adviser Team or its Data Protection: Financial Services Team.