Aona

Generative AI supervision map

Broker-dealer compliance supervisors, technology risk and security teams

Read the complete guide on aona.ai

FINRA’s technology-neutral rules continue to apply when member firms use generative AI, including third-party features. Evaluate the actual use before deployment and connect it to a reasonably designed supervisory system. Input protection, output review and business-record handling are separate responsibilities.

Three fictional member-firm uses illustrate distinct pre-use evaluation and supervisory evidence. Roles are examples, not actual appointments.

Synthetic firm, uses and procedures. No FINRA approval, executed test, communication release or supervisor sign-off is represented.

Your review

Reviewer: ______________________ Date: ______________________

Organisation or team: ____

☐ **Name the use and output destination**

Separate internal assistance, customer communications and use within supervision.

☐ **Define a meaningful test**

Choose a source-backed question and an intentional failure case relevant to that use.

☐ **Assign review and escalation**

Record the responsible roles and actual evidence; do not treat an AI feature as the supervisor.

Checked items record your review, not a compliance approval or an installed-product test result.

Notes and next actions

Decision and scope: ____

Evidence to collect: ____

Owner and next review date: ____

Review matrix

Use	Pre-use check	Supervision and escalation
Internal policy lookup	Compare an answer with the current procedure and its source.	Policy owner resolves contradictions; staff do not rely on unsupported answers.
Draft customer email	Identify factual claims, audience and the applicable communications process.	Assigned reviewer applies the required procedure before any real use or sending.

Use	Pre-use check	Supervision and escalation
AI surveillance summary	Assess source completeness, missed issues, accuracy and limitations.	Supervisory owner retains substantive review and escalates material findings.
Sensitive input in any use	Check permitted data and the exact supported input path.	Security investigates exposure; a passed input check is not approval of the output.
Feature or model change	Identify what changed and which evaluations may be affected.	Use owner reassesses risk and updates procedures where necessary.
Records handoff	Classify the final output and material review evidence.	Records owner assigns applicable retention and archive requirements.

The review follows the use

- **Define:** Task, user, data and destination
- **Evaluate:** Relevant correct and failure cases
- **Supervise:** Named procedure and responsible role
- **Retain:** Required record and review evidence

Safe evaluation cases

1. Internal lookup: Ask an assistant to explain an invented policy rule, then compare the answer with the exact supplied source.
2. Draft communication: Give an invented product description with no performance promise and check whether the output adds one.
3. Supervisory summary: Supply a fictional record set with one intentionally inconsistent entry; check whether the summary preserves it for human review.

These are planned tests, not observed results. No real investor, customer communication or surveillance record is used.

Sources and scope

Sources checked 2026-09-21. Review the current requirements and your actual agreement before using this workbook for a real decision. General information, not professional approval or a completed control test. It does not approve FINRA communications, act as the firm’s supervisor or validate the accuracy of AI outputs.

FINRA Regulatory Notice 24-09. Technology-neutral duties, pre-deployment evaluation, supervision and communications standards for generative AI use.

FINRA Rule 3110: Supervision. Reasonably designed supervisory systems, written procedures and assigned supervisory responsibilities.