

AI and agent adoption roadmap for investment firms

A board-ready plan for CIOs and leadership teams: start with three workflows, build the control layer once, and scale by gate.

```
roadmap.sh

$ grafolio roadmap --plan

phase 1 days 0-90 3 pilots
gate G1 day 90 go/iterate/stop
gate G2 month 6 wave 2
gate G3 month 12 5+ live
gate G4 month 18 level 3-4

✓ control layer: built once
```


01 · THE SITUATION

Most firms have tried AI. Few run it in production they can defend.

The near-term value is operational: efficiency, coverage and speed, not alpha. Data quality and access, not model capability, is the binding constraint.

SOURCE: MERCER 2026 AI IN ASSET MANAGEMENT SURVEY (SUMMARY). DIRECTIONAL.





```
pilot_status.yaml
```

```
workflow: ddq_drafting
```

```
stage: pilot
```

```
model: ok
```

```
blocked_by:
```

```
- process_map: missing
```

```
- data_definitions: missing
```

```
- controls: missing
```

```
- audit_trail: missing
```

02 · THE REAL GAP

The gap between pilot and production is **rarely the model.**

It is the missing process map, the missing data definitions, the missing controls, and the missing link between an AI answer and the source behind it.

03 · WHY NOW

The old template no longer covers the new risk

There is no prescriptive regulatory template for GenAI or agents. The firm must write its own policy, before examiners ask.

FEDERAL RESERVE · APR 2026

SR 26-2

Replaced SR 11-7, the template many firms borrowed. Places generative and agentic AI outside its scope, so firms must design those controls themselves.

FINRA · 2026 OVERSIGHT REPORT

FINRA 2026

Store prompt and output logs, track model versions and keep human-in-the-loop review. For agents: track actions and limit what they can do.

SEC · DIVISION OF EXAMINATIONS

FY2026 exam priorities

Reviews the accuracy of AI claims and whether policies are adequate to supervise AI use, including in back-office and trading functions.

STANDARDS · VOLUNTARY AND CERTIFIABLE

NIST AI RMF · ISO/IEC 42001

The practical starting point: Govern, Map, Measure, Manage, the GenAI Profile (NIST AI 600-1), and a certifiable AI management system.

04 · WHY NOW

Agents raise the value and the control burden together

CHAT ASSISTANT

Answers a question.

A person reads the answer and decides what to do with it.



agent · takes actions

```
> open borrower_financials.pdf
> query loan_system --deal 114
> run covenant_model_v12.xlsx
> write covenant_results.draft
? send needs human approval
```

FINRA 2026 · AGENT-SPECIFIC RISKS

Autonomy without human validation · acting beyond scope or authority · auditability of multi-step reasoning · data sensitivity · gaps in domain knowledge · misaligned reward functions

These are control design problems. They are solvable, if the firm designs for them before deployment.

05 · WHERE TO START

Start with three workflows, not thirty

Choose high-volume, document-heavy, reversible tasks where a qualified person already reviews the output.

$PRIORITY = V \times F \times (6 - R)$

Value and feasibility average 1-to-5 sub-scores. Risk takes the highest, because one severe risk is enough. R = 5 is never a first deployment.

score_candidates.py

#	candidate	V	F	R	score	
1	DDQ / RFP drafting	4.0	4.25	3	51.0	✓
2	Perf. commentary	3.5	4.25	3	44.6	✓
3	Covenant monitoring	4.25	3.5	3	44.6	✓
10	Research synthesis	3.25	3.5	4	22.8	✗
12	Autonomous orders	3.0	2.0	5	6.0	✗

gates: R<=3 · reviewed today · baseline

selected: 1, 2, 3

ILLUSTRATIVE SCORES · HYPOTHETICAL FIRM · ROADMAP \$6.3

06 · THE CONTROL LAYER

Build the control layer once, before go-live

01 Logs

Every prompt, output and tool call is stored.

02 Versions

Model and prompt version recorded for each output.

03 Citations

Every figure traced to its source cell or page.

```
audit_record.json

{
  "inputs": [docs, data]
  "sources": [page, cell]
  "prompt_v": "<id>"
  "model_v": "<id>"
  "tools": [calls...]
  "reviewer": "<name>"
  "decision": "approved"
}
```

04 Least privilege

Agents get only the permissions the step needs.

05 Human checkpoints

Named reviewers at defined points.

06 Exception queue

Anything outside limits routes to a person.

“Show me what the model saw, what it said, and which version said it.”

07 · AUTONOMY AND RISK

Autonomy is set per workflow step, not firm-wide

First deployments start at A1 or A2. Moving up one level takes a full quarter at the current level, evaluation above threshold, a review of edit rates and exceptions, and council approval.

NO AUTONOMOUS TRADING, CLIENT ADVICE OR EXTERNAL COMMUNICATIONS.

A0	Human only	Investment decisions, sign-offs, client advice
A1	Assist	Answers, retrieves and summarizes on request
A2	Draft and recommend	Full draft with citations; human approves each
A3	Act with checkpoints	Defined actions within limits; pauses for approval
A4	Autonomous within bounds	End to end within hard limits; reviewed after

RISK TIER · SET BY THE HIGHEST-RISK STEP

T1 Low Internal productivity; no external output	T2 Moderate Reviewed drafts; confidential data	T3 High Decisions, valuations, filings, MNPI	T4 Restricted Autonomous trading, advice, external comms
---	---	---	---

DATA FOUNDATIONS

Per workflow, not
per enterprise.

69%

of managers cite data constraints as a significant
barrier (Mercer 2026)

08 · DATA

Fix the data the first workflows touch

- | | | |
|----|--------------------------------|---|
| 01 | Metric library | One definition, owner and calculation for each metric in scope. |
| 02 | Golden sources | One authoritative source for each data element the workflow uses. |
| 03 | Governed ingestion | In-scope documents ingested with lineage, treated as untrusted input. |
| 04 | Registered Excel models | Versioned and locked. Agents run them as tools; they never edit formulas. |


09 · BUY, BUILD OR PARTNER

Buy the commodity, build the differentiator, partner for speed



Commodity tasks

Drafting, summarizing, meeting notes and coding help through enterprise licenses with sound data terms, plus vendor AI features where the vendor owns the workflow.



The firm's edge

Configure or build where the workflow crosses documents, Excel and several systems: underwriting, IC memos, covenant monitoring, LP reporting.



Speed and skill

Specialists map processes, stand up the control layer and deliver first deployments, with explicit knowledge transfer.

OWN THE CONTROLS IN EVERY CASE

Logging, versioning, permissions, inventory and evaluation stay under the firm's control. Regulators hold the firm, not the vendor, responsible.

10 · FIRST 90 DAYS

The first 90 days end in a go, iterate or stop decision

GATE 1

DAYS 0-30

Mobilize and assess

- Board approves risk appetite and Phase 1 funding
- Charter the AI council; issue an interim policy
- Build the AI inventory, including vendor AI
- Score candidates; select three workflows

DAYS 31-60

Design and build

- Map the processes; measure baselines
- Build evaluation sets of 25 to 50 cases
- Stand up control layer v1
- Draft the AI policy suite


DAYS 61-90

Pilot and decide

- Run pilots on live data, in parallel
- Test security, including prompt injection
- Deliver the first board AI report
- **Gate 1: go, iterate or stop**

11 · 18-MONTH GATES

Scale is earned at four gates over 18 months



Baselines before deployment

Hours, cycle time, error and rework rates, review time and exceptions, measured first.

LEADING INDICATORS

Adoption · evaluation accuracy · reviewer edit rates · exceptions

LAGGING INDICATORS

Hours released and redeployed · cycle time · post-release errors · incidents · realized value vs plan

Four decisions for the board

- 01

AI risk appetite

AI may assist and prepare; humans stay accountable for decisions, client communications and filings.
- 02

AI council charter

COO or CIO chair, CCO and CRO voting, quarterly reporting.
- 03

Phase 1 funding

Days 0 to 90 now; later phases released only at gates.
- 04

Selection method

Endorse the scoring method; the council picks the first three.

What the council runs

- 01

AI council: COO or CIO chair; CCO and CRO voting
- 02

A RACI for the program
- 03


One AI inventory, including AI inside vendor products
- 04

A written human-in-the-loop policy
- 05

Vendor AI due diligence and contract terms
- 06

MNPI segregation of indexes and agent memory
- 07

Compliance review of every AI claim the firm makes



```
mail
```

```
$ mail contact@grafolio.ai
```

```
--subject "Roadmap"
```

```
# the detailed roadmap covers:
```

- ✓ 18-month plan and gates
- ✓ scoring rubric, worked example
- ✓ control checklist mapped to
NIST, ISO 42001, FINRA, SEC
- ✓ board reporting template

13 · NEXT STEP

Get the detailed roadmap

Email us with **Roadmap** in the subject line.

contact@grafolio.ai

SUBJECT: ROADMAP

GRAFOLIO.AI

Grafolio provides technology and process advisory, not investment, legal or regulatory advice. Workflow examples are illustrative; no client results are implied.