harmonic

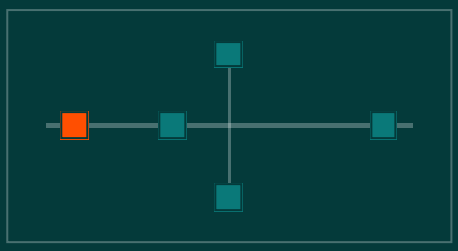

INVESTMENT INDUSTRY GUIDE

The Investment Firm's Guide to Governing AI Adoption

> A practical guide for approving AI faster while protecting deal, investor, fund, portfolio-company, firm, and employee data

Industry Guide | August 2026

START HERE

Turn pressure to adopt AI into a safe route to yes

Investment teams want the latest AI because it can compress research, diligence, memo drafting, analysis, coding, and operating work. Tools such as Claude Cowork, ChatGPT, Copilot, Gemini, and connected agents are becoming more capable quickly. Waiting for a final policy does not stop adoption. It often moves the work into personal accounts and unreviewed workflows.

This guide is a working playbook for security, IT, AI and data, compliance, legal, risk, and operations leaders. Use it to give teams approved ways to work while keeping sensitive investment information inside clear boundaries.

Set the baseline

Decide which tasks, data, tools, accounts, and actions are permitted, conditional, or prohibited.

Approve faster

Reuse decisions by workflow so teams do not wait for a fresh review every time they use AI.

Control actual use

Place guidance and enforcement where people paste, upload, retrieve, connect, or act.

Show value and evidence

Measure adoption, find useful workflows, investigate incidents, and demonstrate oversight.

How to use it: Page 3 shows what the program should make easier. Pages 4-6 establish policy, ownership, and approval. Pages 7-11 cover enforcement, evidence, regulation, incidents, and agents. Page 12 provides a 90-day plan. Page 13 identifies the tools and components worth considering.

The test for a good control is practical: it should protect the firm and make the approved path easier to use.

OPERATING PAYOFF

What gets easier when the controls work

The controls in this guide should reduce friction for employees and for the teams responsible for AI. If they only add review steps, the program will be bypassed. The practical payoff should be visible in everyday work.

Control	What changes for the firm
Live AI inventory	Security and IT stop relying on surveys. They can see which tools and account types teams actually use, then focus reviews and support where demand is real.
Approved-use register	Employees get a quick answer on which tool to use for research, diligence, reporting, coding, or portfolio work. Repeat questions and one-off approvals decline.
Workflow-based decisions	Compliance, legal, security, and data owners approve a reusable pattern instead of reviewing every prompt. Similar work can move faster under the same conditions.
Contextual guidance	People receive a useful warning or approved alternative at the moment of risk. The firm can reserve blocking for clear, high-impact violations.
Scoped evidence and incident data	Investigations start with the tool, account, user, workflow, data category, and action already identified. Audit and incident response take less manual reconstruction.
Usage intelligence	Leaders can see which use cases create value, where training is needed, which approved tools are underused, and where an unapproved tool fills a real capability gap.
Agent boundaries	Teams can pilot Cowork, coding agents, MCP, and connected workflows with narrow permissions and human checkpoints, then expand based on evidence.

Where Harmonic fits

Harmonic is a Workforce AI Security platform, also described as AI Governance and Control. It discovers supported web, desktop, CLI, embedded, and agentic AI use, builds an AI asset inventory, analyzes sensitive information semantically, applies appropriate coaching or controls, and turns the same evidence into Usage Intelligence.

The firm remains responsible for its regulatory, contractual, privacy, and business requirements. Harmonic applies firm policy across supported AI activity; coverage depends on the components and integrations deployed.

POLICY BASELINE

Decide what can go where

Investment information is often sensitive because of context. A target name, valuation assumption, LP detail, or portfolio-company problem may not match a conventional data loss prevention rule, but it can still reveal a strategy, transaction, or material nonpublic information.

Adopt a firm-wide baseline, then apply stricter rules where regulation, contracts, fiduciary duties, information barriers, or business risk require them.

Permitted | Review or conditions | Stop or escalate

Data and use	Typical treatment	Minimum conditions
Public information used for research, brainstorming, or summarization	**Generally permitted**	Firm-approved tool or account; human review when the output informs an investment or external communication.
Internal firm information, including policies, budgets, code, and operational material	**Limited**	Approved enterprise environment, business purpose, and appropriate access, repository, and retention controls.
Confidential target, diligence, valuation, investment-committee, or portfolio-company information	**High control**	Tool approved for the data and workflow, firm-managed account, contractual safeguards, and no conflicting deal restriction.
Investor or LP PII, nonpublic fund information, or unpublished performance	**Restricted**	Specifically approved environment, scoped access, privacy review, and an activity record appropriate to the use.
MNPI, credentials, secrets, production data, or other highly restricted information	**Restricted or prohibited**	Compliance, legal, security, and data-owner review; additional safeguards or de-identification where appropriate.
Sensitive information in a personal or public AI account	**Prohibited**	Block or redirect to an approved enterprise alternative.
External communication, trade or order submission, deletion, transfer, or system change by an agent	**Prohibited by default**	Enable only after specific approval, scoped permissions, logging, action limits, and meaningful human review.

De-identification can reduce risk, but a collection of facts may still identify a fund, investor, target, transaction, or portfolio company.

GOVERNANCE

Establish ownership and visibility

Give one cross-functional group authority to approve tools and uses, resolve exceptions, and change controls. Include security, IT, AI or data leadership, compliance, risk, privacy, legal, procurement, operations, and representatives from investment and investor-facing teams. Name an executive sponsor and an operating owner.

The group should maintain six working deliverables:

Deliverable	Minimum content
AI inventory	Tools, embedded features, accounts, owners, users, teams, surfaces, workflows, connected systems, and review status.
Approved-use register	Approved account, permitted tasks and data, prohibited uses, owner, and next review date.
Data and workflow matrix	Research, diligence, reporting, portfolio, coding, and corporate workflows; data owners, destinations, actions, and required treatment.
Vendor assessment standard	Required security, privacy, contractual, identity, administrative, logging, and investigation capabilities.
Use and exception process	Short rules for common work, named approvers, response target, expiry date, and recorded decisions.
Incident and evidence procedure	Reporting route, containment authority, provider contacts, evidence rules, access model, and notification decision owners.

Inventory use across browsers, desktop applications, embedded SaaS features, coding tools, APIs, local agents, MCP servers, and connectors. Track whether each account is enterprise-managed or personal and where users paste, upload, retrieve, or send data. Refresh the inventory as vendors add features and account settings change.

VENDOR AND USE REVIEW

Approve tools, accounts, and workflows

Approve the specific account, configuration, data type, workflow, and AI surface. An enterprise contract does not make every feature suitable for deal, fund, investor, or portfolio-company information.

Review area	Answer before approval
Data use and retention	Are prompts, files, outputs, or feedback used for training or service improvement? How are active data, backups, and deletion handled?
Provider access and processing	Can personnel or subprocessors access content? Where is it processed, and how is customer data separated?
Enterprise controls	Does the service support SSO, SCIM, roles, administrator controls, audit logs, managed accounts, and configurable retention?
Feature and action risk	Can the feature upload files, retain memory, browse, execute code, retrieve from firm systems, or take actions? Can administrators disable or limit those functions?
Contract and incident support	Do terms limit secondary use, protect confidential information, require incident notice, and provide useful logs and support for an investigation?
Change and exit	How are changes to models, terms, connectors, and data practices communicated? What happens to firm data when the account or contract ends?

Publish the approved-use register where employees already seek guidance. Reassess an approval when the provider, model, feature set, contract, retention behavior, or connected systems change. NIST's Generative AI Profile supports ongoing measurement, monitoring, third-party evaluation, and incident planning.[3]

ENFORCEMENT

Put controls where data moves

Focus controls on the moment a user pastes text, uploads a file, invokes an embedded feature, calls an API, or allows AI to retrieve or change information.

Surface	Minimum control set
Browser and desktop AI	Discover use, distinguish enterprise from personal accounts, inspect relevant content and destination, coach approved alternatives, and block clear violations.
Embedded AI in SaaS	Maintain a feature inventory, review the downstream provider, configure administrator settings, and notify users when the data path changes.
Coding tools and APIs	Approve endpoints, use managed credentials, apply repository and data rules, cover IDE and terminal workflows, and log relevant activity.
Agents, MCP, and connectors	Inventory servers and actions, use dedicated identities and least privilege, require approval for high-impact actions, retain audit trails, and support rapid revocation.

Make controls workflow-aware

- Set clear rules for research, origination, diligence, investor reporting, portfolio support, coding, and corporate work.
- Record the sensitive data, approved tools and accounts, action limits, owner, and review date for each workflow.
- Make restrictions visible at the point where an employee selects a tool, account, or action.
- Provide a fast route for clarification, approval, and time-limited exceptions.

Block when the violation is clear or the likely harm is high. When context is uncertain, give a short warning that names the issue and points to an approved option or exception path.

EVIDENCE AND PRIVACY

Create evidence without turning governance into surveillance

Regulatory obligations vary by entity and activity. For covered institutions, the 2024 amendments to Regulation S-P include written incident-response and recordkeeping requirements.[1] The SEC's 2026 examination priorities include compliance with those amendments.[2]

Governance need	Operating requirement
Applicable obligations	Map the entity, product, jurisdiction, client and investor commitments, information barriers, recordkeeping rules, and privacy requirements that apply to each workflow.
Attribution and activity records	Capture the user or responsible owner, tool, account, surface, workflow, data category, policy decision, action, and outcome where supported.
Review and access	Use role-based access, masking, exclusions, and scoped investigations. Limit raw prompt, file, and output access to people with a defined need.
Retention and response	Set retention, deletion, export, legal hold, and incident-evidence rules before collecting detailed activity data.
Adoption and value	Report active use, workflow mix, approved-tool share, coaching outcomes, and license signals at team level where possible.

Set monitoring guardrails

- Decide who can see tool names, user identities, summaries, alerts, raw prompts, files, and outputs before collection begins.
- Limit raw content access to defined investigations and audit that access.
- Publish the purpose of monitoring, the access model, and the employee exception route.
- Report adoption, use cases, policy events, and unresolved exceptions without routinely exposing sensitive content.

The activity record should show how the firm governed AI without creating a broadly accessible second copy of sensitive investment work.

REGULATORY MAP

Map AI use to the rules that already govern the firm

Existing requirements covering customer information, MNPI, records, communications, supervision, and compliance programs can apply when AI enters a regulated workflow. The SEC's 2026 examination priorities specifically address automated investment tools and AI technologies, as well as Regulations S-P and S-ID where applicable.[2]

Regulatory area	Who should assess it	What AI governance should evidence
Compliance programs **Rules 206(4)-7 and 38a-1**	Registered advisers and investment companies	Owners, inventory, risk assessment, written policy, approvals, testing, exceptions, and review inputs.
Regulation S-P	Covered broker-dealers, investment companies, RIAs, and transfer agents	Customer-information safeguards, provider oversight, incident response, evidence, and notification decisions.
MNPI and codes of ethics	Investment advisers and personnel covered by firm policy	Information barriers, contextual controls, approved destinations, access limits, escalation, and guidance.
Books and records **Rules 204-2, 31a-1, and 31a-2**	Registered advisers and investment companies	Required AI-assisted records, approved channels, retrieval, retention, and legal hold.
Marketing and anti-fraud	Advisers and firms making investor-facing representations	Substantiation, fair balance, human approval, required records, and controls against misleading AI claims.
Regulation S-ID	Covered firms with covered accounts	Identity-theft red flags, response, training, program updates, and provider oversight.

Conditional overlays: Assess FINRA and Regulation Best Interest for broker-dealer affiliates, New York DFS Part 500 and other state privacy or cybersecurity rules for covered entities, and relevant international requirements. Scope applicability entity by entity and workflow by workflow.

INCIDENT RESPONSE

Prepare for AI-related data incidents

Employees should report a mistaken upload, unsafe connector, or unintended AI action immediately, without first deciding whether it qualifies as a breach.

Phase	Required action	Primary owners
Contain	Stop use, revoke a connector, token, or credential if needed, prevent further sharing or action, and preserve the minimum necessary evidence.	Security and IT
Establish scope	Identify the tool, account, user, surface, workflow, data, fund, investor, target, portfolio company, time, recipients, and provider retention terms.	Security, risk, compliance, and data owner
Seek provider action	Request deletion or restricted processing, confirm retention behavior, preserve logs, and use contractual incident channels.	Procurement, legal, privacy, and security
Assess obligations	Decide client or investor notice, contract, privacy, regulatory, insurance, information-barrier, and privilege obligations.	Legal, compliance, privacy, risk, and business owner
Correct	Update controls, tool status, guidance, training, workflow rules, or access and assign follow-up actions.	AI governance group

Do not promise deletion until the provider confirms it. Do not delay containment while debating the formal definition of the event. Legal and compliance should decide notification, privilege, and regulatory issues.

Run a tabletop using a realistic scenario, such as a valuation model uploaded through a personal account or an agent retrieving information from the wrong fund or portfolio company. Test the contact list, provider escalation path, decision authority, and available evidence.

AGENT CONTROLS

Set boundaries before enabling agents

Agents can retrieve data and act across research platforms, data rooms, email, collaboration, CRM, code, and portfolio systems. The risk changes when a system can choose tools, move data, and take action without a separate prompt for every step.

Before deployment

- [ ] Start with a narrow workflow, low-sensitivity data, and a named human checkpoint.
- [ ] Use a dedicated identity and least privilege. Scope credentials to the user, system, task, fund, deal, or portfolio company, and use read-only access where possible.
- [ ] Inventory and approve MCP servers, connectors, tools, permissions, and external providers. Deny unknown or unnecessary connections by default.
- [ ] Require approval before external communication, trade or order submission, record changes, deletion, transfers, or other high-impact actions.
- [ ] Test indirect prompt injection, hidden instructions, retrieval boundaries, cross-deal access, action limits, and failure behavior.
- [ ] Log retrieval, tool calls, actions, approvals, and the responsible user. Set transaction and time limits and provide a fast way to pause the workflow and revoke access.

Reassess permissions and tests whenever the model, connector, MCP server, or agent behavior changes. Expand access only after testing shows that the system stays within its assigned workflow and data boundaries.

IMPLEMENTATION PLAN

A 90-day implementation plan

Use the first 90 days to establish a working baseline. The CIO or CISO should sponsor the program with AI or data leadership, compliance, privacy, legal, IT, and representatives from investment and operating teams.

Period	Priority	Required outputs	Evidence of completion
Days 1-30	Discover and define policy	Governance group, named owners, AI inventory, account distinctions, priority workflows, sensitive-data map, privacy rules, and initial approved-use register.	Meeting cadence, coverage map, adoption baseline, policy matrix, and decision log.
Days 31-60	Steer use and close coverage gaps	Contextual coaching and intervention, exception route, endpoint and priority coding coverage, role-based training, masking and access rules, and tested incident workflow.	Configured controls, false-positive review, training report, access model, and completed incident exercise.
Days 61-90	Govern agents and report value	Agent and MCP inventory, connector and action rules, leadership reporting, policy tuning, adoption and license review, and portfolio baseline where relevant.	Agent test record, remediation log, usage report, next-quarter roadmap, and reusable rollout kit.

At each checkpoint, review actual use, employee feedback, unresolved exceptions, and controls that repeatedly interrupt permitted work. Adjust the policy and approved options before expanding enforcement.

Selected references

1. U.S. Securities and Exchange Commission, "Regulation S-P: Privacy of Consumer Financial Information and Safeguarding Customer Information," final rule, May 16, 2024: sec.gov/rules-regulations/2024/06/s7-05-23
2. U.S. Securities and Exchange Commission, "Fiscal Year 2026 Examination Priorities," November 2025: sec.gov/files/2026-exam-priorities.pdf
3. National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile," NIST AI 600-1, July 2024: nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

PRACTICAL TOOLS

Tools that materially change the program

You can establish ownership, policy, approval criteria, and an incident path without buying a new platform. The components below become useful when the firm needs live coverage, repeatable evidence, or a safe way to expand into Cowork, coding tools, and connected agents.

Tool or component	What it changes	Reach for it when	Where to start
Harmonic Explore, Guide, and Command	**Explore** adds the AI asset inventory, prompt-level browser and supported OTel visibility, application risk profiles, and Usage Intelligence. **Guide** adds browser and endpoint governance with warn, nudge, or block controls. **Command** adds supported MCP, agent-action, and connector controls.	Use Explore for visibility and evidence, Guide for real-time browser and endpoint guardrails, and Command when AI connects to enterprise systems or teams adopt MCP and agents.	harmonic.security Confirm current package and surface coverage.
Enterprise AI workspaces and identity	Moves users into firm-managed accounts with centralized access, administrator settings, lifecycle management, and vendor terms suitable for business use.	Employees are using personal accounts, or confidential work is moving beyond public research and brainstorming.	Review the enterprise administration, data, retention, SSO, and SCIM documentation for each approved provider and identity platform.
Microsoft Purview DSPM for AI	Adds Microsoft-native discovery, sensitive-data analysis, audit, DLP, retention, eDiscovery, and investigation support for Copilot and other covered AI activity.	The firm runs heavily on Microsoft 365 and wants AI controls tied to existing labels, audit, DLP, and compliance workflows.	learn.microsoft.com/purview/ai-microsoft-purview
Native telemetry and compliance feeds	Exports activity from supported desktop and coding tools through OpenTelemetry, audit, or compliance feeds so it can be retained, investigated, and analyzed with other AI activity.	Cowork, Claude Code, Codex, or another desktop or CLI tool is moving beyond a small pilot and browser controls do not cover it.	Harmonic OpenTelemetry guide and each provider's monitoring documentation.
AI gateway for API and developer traffic	Centralizes authentication, provider routing, usage tracking, budgets, and audit logging for managed model access.	Teams build against model APIs, use coding assistants through firm infrastructure, or need consistent controls across model providers.	Anthropic LLM gateway guide and the firm's cloud or gateway standard.
Agent, MCP, and skill assessment	Inventories local agent configuration and scans third-party skills or extensions for risky permissions, prompt injection, data exfiltration patterns, embedded credentials, and dangerous operations.	Before enabling MCP servers, connectors, third-party skills, Cowork extensions, or autonomous workflows on firm devices.	claudit-sec and Cisco Skill Scanner

Start with the gap, not the product list. Inventory actual use first. Add a component only when it closes a defined coverage, evidence, identity, or response gap.