ALPHA
intelligence
LABS

AI Governance, Policy, and Operating Guide for Investment Offices

A Living Institutional Policy Document

Monthly Update Edition · Version 1.1 · Effective July 9, 2026

Contents

Introduction & How to Use This Document

Monthly Update & Change Log

1. Purpose & Scope

2. AI Philosophy & Fiduciary Principles

3. Governance Structure & Accountability

4. Risk Classification Framework

5. Human-in-the-Loop Controls

6. Approved Technology & Vendor Governance

7. Data Governance, Privacy & Security

8. Auditability, Documentation & Monitoring

9. Metrics, Reporting & Operational Alpha

10. Permissible & Prohibited Uses

11. Training, Culture & Responsible Experimentation

12. Policy Review & Change Management

Board Resolution — Adoption of AI Governance Policy

Trustee Cover Memo: AI Governance Policy

Staff AI Playbook — Practical Guidance

Expanded Annexes — Detailed Implementation Guide

Version Control & Change Log

Introduction & How to Use This Document

This document is a starting point: a comprehensive, board-ready AI governance framework that an investment office can adopt and adapt to its own structure and risk tolerance.

> When presenting this to a board• Approve Sections 1–9 as formal policy.• Acknowledge the Annexes as reference material.• Distribute Sections 10–11 separately to staff.That single move avoids two failure modes: over-lawyering staff behavior, and under-governing fiduciary risk.

INTENDED FOR

Trustees · Auditors · Regulators · External counsel

Document Structure & Intended Use

This document is intentionally structured to serve multiple audiences without creating confusion or unnecessary bureaucracy. The first sections are written for boards and executive leadership and are intended to be formally approved and referenced in oversight and audit contexts.

Later sections translate these principles into institutional operating policies and staff-level guidance. Annexes provide detailed implementation tools and reference frameworks that support compliance and audit readiness, without burdening day-to-day users with excessive detail.

Monthly Update & Change Log

Date of Update: July 9, 2026 · Version 1.1

High-Level Summary of Changes

This update incorporates recent AI governance developments relevant to institutional investment offices, including board-level AI oversight expectations, revised model-risk-management principles, increased attention to agentic AI and autonomous workflows, enhanced third-party and vendor-governance expectations, and emerging transparency, logging, and documentation obligations for AI systems.

Why These Changes Matter for an Investment Office

Investment offices are increasingly using AI to support research, manager monitoring, document review, reporting, workflow automation, and decision-informing analytics. Recent regulatory and peer-institution developments emphasize that AI governance must be board-defensible, auditable, risk-based, and clear about human accountability. For resource-constrained investment offices, the practical implication is not to create a large new bureaucracy, but to maintain a clear inventory, assign human ownership, document material uses, control confidential data, validate outputs, and escalate higher-risk AI activity before it affects fiduciary decisions.

Sections Materially Updated

Sections 1, 3, 4, 5, 6, 7, 8, 10, 11, 12, Annex B, Annex C, Annex D, Annex E, and Annex F.

Sources Reflected in This Update

This update reflects recent developments including board-level AI governance work, revised model-risk-management guidance, EU AI Act implementation materials, transparency expectations for AI-generated content, and recent financial-services regulatory focus on agentic AI, concentration risk, and reliance on large technology providers.

1. Purpose & Scope

The purpose of this policy is to define how AI may be evaluated, approved, deployed, used, monitored, modified, and retired within the investment office. AI is defined broadly to include machine learning models, generative AI systems, large language models, retrieval-augmented systems, AI agents, automation tools, model-based analytics, vendor-provided AI functionality, and advanced analytics capable of influencing workflows, communications, risk assessments, investment analysis, operational processes, or decision-making.

This policy applies to all investment, operations, compliance, finance, and administrative functions, as well as third-party vendors acting on behalf of the organization. All AI usage must align with fiduciary duties owed to beneficiaries and comply with applicable legal, regulatory, and ethical standards.

For avoidance of doubt, this policy applies whether AI is developed internally, embedded in third-party software, provided by an investment manager, offered by a consultant, accessed through a public tool, or used as part of a workflow automation or agentic system. AI systems that do not directly make investment decisions may still be subject to this policy if they influence research, diligence, reporting, compliance, operational controls, external communications, or the evidentiary record supporting fiduciary decisions.

2. AI Philosophy & Fiduciary Principles

AI is a tool, not a decision-maker. Its role is to augment human expertise by improving efficiency, surfacing insights, and reducing operational friction, while preserving human judgment as the final authority for all fiduciary decisions.

The organization commits to using AI in a manner that is transparent, explainable, and accountable. All material AI outputs must be understandable to oversight bodies, and responsibility for outcomes remains with designated human owners at all times.

3. Governance Structure & Accountability

The Board of Trustees retains ultimate oversight responsibility for AI governance. Management, led by the CIO and COO, is accountable for ensuring that AI usage aligns with strategic objectives, fiduciary standards, and risk tolerances.

An AI Steering Committee oversees implementation, approves tools and use cases, assigns risk classifications, and escalates material issues. Clear ownership is assigned to each AI system to ensure accountability throughout its lifecycle.

The Board's role is oversight, not day-to-day technology approval. Board oversight should focus on whether management has established an appropriate AI governance framework, risk appetite, escalation process, reporting cadence, and control environment. Management remains responsible for implementation, tool selection, use-case approval, staffing, monitoring, and remediation.

The AI Steering Committee shall maintain an inventory of approved AI systems and material use cases. The inventory should identify the business owner, purpose, risk classification, data types used, vendor or internal platform, human oversight points, review cadence, and whether the system produces decision-informing, externally distributed, or beneficiary-impacting outputs.

Management shall report to the Board at least annually, and more frequently where material issues arise, regarding AI adoption, material use cases, control effectiveness, incidents, vendor exposure, regulatory developments, and any changes to risk appetite or governance structure.

4. Risk Classification Framework

AI use cases are classified based on their potential impact on investment outcomes, beneficiaries, confidential information, operational resilience, regulatory exposure, external communications, and the organization's fiduciary record. Risk classification determines the level of review, approval, documentation, validation, monitoring, and escalation required.

Low-risk use cases generally include drafting, summarization, formatting, translation, workflow assistance, and administrative support using approved tools and non-confidential or properly protected information. Medium-risk use cases include investment research support, manager screening assistance, exposure analysis, portfolio reporting support, operational exception review, compliance support, or analytics that may inform but do not determine fiduciary decisions. High-risk use cases include decision-informing models, autonomous or semi-autonomous agents, systems that affect external reporting, systems that process sensitive personal or confidential information at scale, systems embedded in investment or risk models, and systems that may materially affect beneficiaries, counterparties, managers, or the organization's fiduciary record.

AI systems may not be classified solely by the underlying technology. Classification must consider the use case, data sensitivity, degree of automation, reliance by staff, potential harm if outputs are wrong, reversibility of actions, vendor dependency, and whether outputs are used in board materials, investment recommendations, manager decisions, compliance determinations, or external communications.

The AI Steering Committee may reclassify any AI system or use case if actual usage, data inputs, autonomy, integration, regulatory expectations, or risk profile changes over time.

5. Human-in-the-Loop Controls

Human oversight is mandatory for all AI-assisted activities that influence decisions or external communications. AI outputs must be reviewed by qualified personnel before use, and no AI system may execute autonomous fiduciary actions.

Human-in-the-loop checkpoints are documented for medium- and high-risk use cases, ensuring that accountability is explicit and auditable. Responsibility for outcomes cannot be delegated to AI systems or vendors.

For medium- and high-risk AI uses, human review must be substantive rather than ceremonial. Reviewers must have sufficient subject-matter expertise to evaluate the reasonableness of the output, identify missing context, challenge assumptions, and determine whether the output may be used, revised, escalated, or rejected.

Agentic AI systems require heightened controls. No AI agent may initiate, approve, transmit, modify, delete, or execute a material investment, operational, compliance, financial, contractual, or external communication action unless the action has been specifically approved within the system's authorized workflow and remains subject to documented human oversight. All agentic workflows must identify what systems the agent may access, what actions it may take, what actions are prohibited, what approvals are required, and how actions can be reviewed or reversed.

6. Approved Technology & Vendor Governance

Only AI tools and platforms approved by management may be used. Approval considers data protection, security posture, model transparency, and alignment with fiduciary objectives. Enterprise-grade controls are required for all approved tools.

Vendors, consultants, investment managers, and technology providers using AI on behalf of the organization must demonstrate appropriate security, confidentiality, data-governance, model-risk, auditability, and responsible-AI controls. Contracts should include audit rights, data-use restrictions, confidentiality obligations, incident notification requirements, subcontractor disclosure, data-retention and deletion obligations, cooperation with regulatory or audit inquiries, and obligations to support responsible and explainable AI use.

Vendor due diligence should address whether organizational data, portfolio information, manager materials, beneficiary information, or confidential documents are used to train, fine-tune, improve, or benchmark vendor models. Unless expressly approved in writing, confidential investment-office data may not be used for vendor model training or shared across customers. Vendor AI tools must provide sufficient logging, access controls, documentation, and output traceability to support audit readiness and fiduciary oversight.

7. Data Governance, Privacy & Security

Data integrity and confidentiality are foundational to trustworthy AI. All data used by AI systems must be lawfully obtained, appropriately classified, and protected through encryption, access controls, and monitoring.

Sensitive, personal, or confidential information may only be used within approved environments and in accordance with retention and redaction policies. Data governance practices are reviewed regularly to reflect evolving risks.

The organization shall maintain controls to prevent confidential, proprietary, personal, or legally restricted information from being entered into unapproved public AI systems. This includes portfolio holdings, manager diligence materials, investment committee materials, beneficiary information, employee information, legal advice, nonpublic financial information, credentials, contractual terms, and documents subject to confidentiality or nondisclosure obligations.

AI systems using internal data must be subject to access controls, least-privilege permissions, retention rules, redaction standards, and logging appropriate to the sensitivity of the data. Where feasible, AI systems should preserve source references, retrieval history, user prompts, material outputs, and review status sufficient to support audit, supervision, and reconstruction of material decisions.

8. Auditability, Documentation & Monitoring

Each AI system must be documented with a clear purpose, business owner, technical or vendor owner, risk classification, approved users, data sources, permitted uses, prohibited uses, human oversight points, validation approach, limitations, monitoring cadence, and retirement or suspension criteria. Documentation should be sufficient to explain how the system is used, what it is intended to support, what it is not permitted to do, how outputs are validated, and how material issues are escalated.

For medium- and high-risk AI systems, the organization shall maintain a model or system card, implementation checklist, approval record, testing evidence, and monitoring record. Documentation must be understandable to non-technical stakeholders, including trustees, auditors, regulators, external counsel, and senior management.

Monitoring processes are established to detect errors, drift, misuse, unauthorized access, inappropriate reliance, data leakage, vendor changes, performance degradation, and control failures. Material issues are escalated promptly to management and, where appropriate, to the Board.

AI systems that materially change in functionality, autonomy, data access, vendor configuration, model version, integration, or intended use must be reassessed before continued use. The absence of a material incident does not eliminate the need for periodic review.

9. Metrics, Reporting & Operational Alpha

AI adoption is evaluated using both quantitative and qualitative metrics, including time saved, process improvements, and insights generated. These metrics support accountability and continuous improvement.

Where appropriate, management may translate these measures into estimates of operational alpha or investment impact for reporting to leadership and trustees.

10. Permissible & Prohibited Uses

STAFF GUIDANCE

Permissible uses of AI include drafting, summarizing, research assistance, and workflow acceleration within approved tools. These uses are intended to enhance productivity while maintaining quality and oversight.

Prohibited uses include uploading confidential data into public or unapproved AI tools; relying on AI outputs without verification; representing AI-generated output as final institutional judgment without review; using AI to make autonomous fiduciary, investment, legal, compliance, hiring, or beneficiary-impacting decisions; using unapproved AI agents; bypassing access controls; using AI to create misleading, deceptive, or unverifiable content; or deploying tools that have not been reviewed through the organization's approval process.

Staff must clearly distinguish between AI-assisted drafting and approved institutional positions. AI may assist in preparing analysis, summaries, and preliminary materials, but final conclusions, recommendations, manager views, portfolio actions, board materials, and external communications must reflect human review and accountable institutional judgment.

11. Training, Culture & Responsible Experimentation

The organization encourages responsible experimentation with AI to improve outcomes and efficiency. Training programs ensure staff understand risks, limitations, and ethical considerations.

Annual training on AI governance, data protection, and responsible use is mandatory for all relevant staff.

Training shall include practical instruction on prompt discipline, verification of AI outputs, confidentiality restrictions, hallucination risk, source citation, use of approved tools, escalation of uncertain outputs, and the distinction between productivity use and decision-informing use. Staff responsible for medium- or high-risk AI systems shall receive role-specific training on documentation, monitoring, vendor controls, and human review obligations.
The organization encourages responsible experimentation within approved environments. Experimentation outside approved tools or using confidential information in unapproved environments is prohibited.

12. Policy Review & Change Management

This policy is a living document reviewed quarterly by management and annually by the Board. Updates reflect changes in technology, regulation, peer practices, vendor capabilities, organizational use cases, incident history, and risk tolerance.

Material changes to AI usage, including deployment of agentic systems, use of sensitive data, integration into investment or compliance workflows, external use of AI-generated content, or vendor changes that alter data access or model behavior, require reassessment under this policy.

All revisions are version-controlled, dated, documented, and communicated to staff and relevant stakeholders. Material policy changes require Board review and approval. Non-material operating updates may be approved by management and reported to the Board through the regular AI governance reporting process.

Board Resolution — Adoption of AI Governance Policy

WHEREAS, the Board of Trustees recognizes the growing role of Artificial Intelligence (AI) in investment and operational processes; and

WHEREAS, the Board has a fiduciary obligation to ensure that AI is governed prudently, transparently, and in alignment with beneficiary interests;

RESOLVED, that the Board hereby adopts the Comprehensive AI Governance, Policy, and Operating Guide for Investment Offices, and directs management to implement, maintain, and periodically review the policy.

RESOLVED FURTHER, that management shall maintain an AI governance inventory, risk classification process, human oversight framework, vendor-governance process, and documentation standard sufficient to support fiduciary oversight, audit readiness, and responsible AI adoption.

RESOLVED FURTHER, that management shall report at least annually to the Board on AI usage, risks, and governance effectiveness.

Trustee Cover Memo: AI Governance Policy

This memo accompanies the Comprehensive AI Governance Policy for Board approval. The policy is designed to enable responsible AI adoption while preserving fiduciary accountability. It reflects best practices from peer institutions and leading regulatory frameworks.

The current update reflects the market's movement from general AI experimentation toward governed deployment of AI systems, including generative AI, model-based analytics, and agentic workflows. The policy continues to preserve human fiduciary accountability while giving management practical tools to evaluate, approve, monitor, and retire AI systems in a resource-constrained investment office.

The Board is asked to approve Sections 1–9 as formal policy, acknowledge the Annexes as reference material, and delegate staff guidance to management. This structure balances innovation, oversight, and practicality.

Staff AI Playbook — Practical Guidance

This playbook provides practical guidance for using AI responsibly within the investment office. AI is intended to support your work, not replace judgment.

You may use approved AI tools for drafting, summarization, research assistance, document review, workflow acceleration, and other approved productivity uses. You must not upload confidential data into public tools, rely on AI without verification, use unapproved systems, allow AI to make decisions on your behalf, or use AI-generated content externally without appropriate review.

When in doubt, escalate questions to management or the AI Steering Committee.

Expanded Annexes — Detailed Implementation Guide

Annex A — Board-Ready Sample Clauses

This annex provides formal, board-ready clauses that may be adopted verbatim or adapted to reflect the organization's governance structure. These clauses are designed to clearly establish fiduciary accountability, oversight authority, and escalation rights related to AI systems.

The clauses are intentionally conservative in tone and are aligned with emerging regulatory, audit, and litigation expectations. They are suitable for inclusion in board resolutions, committee charters, investment policy statements, or enterprise risk frameworks.

RECOMMENDATIONS

- Adopt these clauses formally via board resolution rather than informal acknowledgment.
- Reference these clauses in committee charters, including Risk, Audit, and Technology.
- Use the termination clause explicitly in vendor contracts involving AI-enabled services.

Annex B — AI Implementation Checklists

These checklists operationalize AI governance by translating policy principles into concrete actions at each stage of the AI lifecycle. They ensure consistency, reduce reliance on individual judgment, and support audit readiness. The checklists are designed to scale — smaller offices may combine steps; larger organizations may formalize each checkpoint within workflow or ticketing systems.

ADDITIONAL CHECKLIST ITEMS FOR 2026 AND BEYOND

- Confirm whether the use case involves generative AI, model-based analytics, or agentic AI.
- Confirm whether the system can access, change, transmit, delete, or act upon institutional data.
- Identify all connected systems, data sources, and output destinations.
- Determine whether outputs may influence investment decisions, manager decisions, compliance actions, board reporting, or external communications.
- Confirm whether confidential, personal, proprietary, or legally restricted data will be used.
- Document the required human review step before production use.
- Confirm whether vendor data is used for model training, benchmarking, or improvement.
- Establish logging, retention, and review requirements.
- Define suspension, rollback, and incident-escalation procedures.
- Reassess the use case after material model, vendor, workflow, or data changes.

Annex C — Model & System Cards

Model and system cards provide a standardized record for every AI system in use. Their purpose is to make AI systems understandable to non-technical stakeholders, including trustees, auditors, and regulators. They document intended use, limitations, data sources, human oversight points, and risk classification — the backbone of AI auditability.

MINIMUM MODEL OR SYSTEM CARD FIELDS

- System name and vendor or internal owner.
- Business owner and accountable executive.
- Approved purpose and prohibited uses.
- Risk classification and rationale.

- User groups and access permissions.
- Data inputs and data-sensitivity classification.
- Whether the system uses generative AI, predictive modeling, retrieval, automation, or agentic functionality.
- Human review points and approval requirements.
- Output types and intended consumers.
- Known limitations, assumptions, and failure modes.
- Monitoring cadence and performance indicators.
- Logging, retention, and audit evidence.
- Vendor model-training and data-use terms.
- Date approved, date last reviewed, and next review date.

Annex D — Generative AI Usage Guide for Staff

This annex establishes clear, practical guardrails for staff use of generative AI tools, enabling productivity gains while minimizing confidentiality, accuracy, and reputational risks. Generative AI is probabilistic and may produce errors or hallucinations; accordingly, verification and human judgment are mandatory.

Staff must treat generative AI outputs as drafts or analytical aids, not as authoritative records. AI-generated summaries, analysis, recommendations, code, calculations, or communications must be checked against reliable source materials before use. Where AI output is used in investment research, diligence, compliance, board materials, or external communications, the reviewer must confirm that the output is accurate, complete, appropriately sourced, and not misleading.

Staff should not use AI-generated citations, quotations, calculations, legal statements, performance figures, manager descriptions, or policy language without verification. Where the AI system provides source references, staff must confirm that the cited source exists and supports the relevant statement.

Annex E — Data Retention, Redaction & Disposal Rules

This annex defines how data associated with AI systems must be retained, protected, anonymized, and ultimately disposed of. It aligns AI data practices with fiduciary recordkeeping obligations and privacy regulations. Clear retention and disposal rules reduce regulatory risk and prevent unnecessary data accumulation.

AI-related records should be retained in a manner consistent with the organization's records-management, fiduciary, audit, and litigation-hold obligations. Depending on risk classification, records may include prompts, source documents, retrieved passages, generated outputs, reviewer notes, approvals, exception logs, vendor notifications, incident records, and system-card updates. Retention rules should distinguish between routine low-risk productivity use and medium- or high-risk decision-informing use; for higher-risk use cases, records should be sufficient to reconstruct the basis for reliance, identify the human reviewer, and demonstrate that the AI system operated within approved parameters.

Annex F — Industry Framework Alignment & Best Practices

This annex maps the organization's AI governance approach to widely recognized frameworks and regulatory expectations, including the NIST AI Risk Management Framework, ISO/IEC 42001, the EU AI Act, revised interagency model-risk-management guidance, fiduciary governance principles, and relevant financial-services AI oversight developments. Alignment demonstrates prudence and due care, even where formal regulatory requirements are still evolving.

The organization recognizes that not all AI governance frameworks apply directly to investment offices. Nevertheless, alignment with recognized frameworks demonstrates prudence, due care, and a defensible control environment, particularly in areas such as inventory management, risk classification, human oversight, documentation, monitoring, vendor governance, transparency, and incident escalation.

Version Control & Change Log

Version 1.0 — Initial Board Adoption. Initial adoption of comprehensive AI governance framework.

Version 1.1 — July 9, 2026 Monthly Update. Updated to reflect recent AI governance developments affecting institutional investment offices, including board-level AI oversight, revised model-risk-management principles, agentic AI controls, vendor and third-party AI governance, data-use restrictions, documentation expectations, transparency and logging considerations, and staff-facing controls for generative AI. Sections materially updated: 1, 3, 4, 5, 6, 7, 8, 10, 11, 12, Annex B, Annex C, Annex D, Annex E, and Annex F.

All future updates will be versioned, dated, and documented with a summary of changes, approval authority, effective date, and sections materially revised. Material changes require Board review and approval.