ICFDT International Council for Derivative Trading

CFIA

Chartered Financial Intelligence Architect

BODY OF KNOWLEDGE

Examination Standards & Curriculum Framework

First Edition

www.icfdt.com

About This Document

The CFIA Body of Knowledge (BOK) is the official reference document defining the scope, structure, and content of the Chartered Financial Intelligence Architect (CFIA) examination. It serves as the authoritative guide for candidates preparing for the exam, and as the foundational syllabus reference for Approved Prep Providers developing educational materials.

This document does not constitute a study guide or textbook. It defines what candidates are expected to know, understand, and be able to apply. The official preparation resource is the CFIA Official Curriculum, published separately by ICFDT.

Contents

Examination Overview

Attribute	Detail
Format	200 multiple-choice questions
Duration	4 hours
Passing Score	Determined by standard-setting process
Delivery	Computer-based; available year-round
Language	English
Recertification	20 hours CPD annually; renewal every 3 years

Examination fees, registration procedures, and scheduling are published on the ICFDT website.

Exam Weight by Module

#	Module	Weight
1	**The AI-Native Investment Firm**	**12%**
2	**Foundation Models & AI Systems**	**13%**
3	**Investment Data Infrastructure**	**16%**
4	**AI for Alpha Generation & Research**	**18%**
5	**Risk, Compliance & AI Governance**	**20%**
6	**Vendor Strategy & Build/Buy Architecture**	**11%**
7	**AI Leadership, Ethics & Fiduciary Duty**	**10%**

Learning Outcome Statement Format

Each topic within the CFIA BOK is accompanied by one or more Learning Outcome Statements (LOS). Each LOS describes what a candidate must be able to do, using action verbs that indicate the required cognitive level:

- **Describe / Explain:** recall and articulate a concept accurately
- **Apply:** use a framework or concept in a given scenario
- **Evaluate / Analyze:** assess, compare, or critique approaches
- **Design / Construct:** build or specify a framework, policy, or architecture

Proprietary Frameworks

The CFIA examination tests six proprietary frameworks developed by ICFDT and embedded throughout the official curriculum. These frameworks are examinable across multiple modules and should be understood at the level of structure, application, and diagnostic use rather than recall alone.

Framework	Description
CFIA AI Maturity Framework™	Five-stage model (AI-Unaware through AI-Native) for assessing organizational AI adoption. Tested in Modules 1, 5, and 7.
CFIA AI Investment Stack™	Five-layer architecture (Data Infrastructure through Portfolio Decision Engine) showing dependency relationships in AI-enabled investment firms. Tested in Modules 1, 3, 4, and 6.
CFIA AI-Augmented Investment Loop™	Six-stage continuous investment cycle (Data → Hypothesis → Simulation → Decision → Feedback → Update). Tested in Modules 1, 3, and 4.
CFIA Data Provenance Chain™	Six-stage data governance standard (Origination → Validation → Transformation → Storage → Retrieval → Deprecation). Tested in Modules 3 and 5.
CFIA Signal Integrity Framework (SIF)™	Five mandatory tests for AI-generated investment signals before deployment (Statistical Validation, Walk-Forward Testing, Regime Analysis, Capacity Analysis, Live Pilot). Tested in Modules 4 and 5.
CFIA AI Governance Trident™	Three-pillar governance structure (Model Risk Management, Explainability & Audit, Regulatory Alignment). Tested in Module 5.

MODULE 1 | Exam Weight: 12% | Estimated Study Time: 6–8 hours

The AI-Native Investment Firm

This module establishes the strategic context for AI adoption in investment management. Candidates must understand what AI can do in an investment setting, and how investment organizations evolve in structure, culture, talent, and process in order to benefit from it sustainably.

1.1 AI Maturity in Investment Organizations

Learning Outcome Statements (LOS)

- Define the AI-native investment firm and distinguish it from AI-enabled or AI-aware organizations.
- Apply the CFIA AI Maturity Framework to assess a firm's current stage and identify gaps.
- Explain the structural and cultural characteristics of firms at each of the five maturity stages.
- Identify the primary risks associated with each maturity stage, including competitive disadvantage, shadow AI adoption, governance failure, and capability illusion.

Required Knowledge

- **The CFIA AI Maturity Framework™, five stages:**
  - Stage 1: AI-Unaware. No AI in production; risk is competitive disadvantage and talent loss
  - Stage 2: AI-Aware. Individual experimentation without governance; risk is shadow AI adoption
  - Stage 3: AI-Enabled. AI deployed in production but additive rather than transformative; risk is integration depth
  - Stage 4: AI-Integrated. AI embedded in core workflows; governance in place; structural inflection point
  - Stage 5: AI-Native. AI is the operating model; proprietary advantage; human judgment correctly allocated
- Stage 3 against Stage 4: the most commonly examined boundary, turning on additive rather than embedded architecture
- Capability Illusion: mistaking tool adoption for organizational AI capability (a Stage 2–3 risk)
- Why AI-nativeness is a structural and cultural state rather than a technology procurement decision

> **CFIA PROPRIETARY FRAMEWORK™ | CFIA AI Maturity Framework™**
>
> *Five-stage diagnostic model from AI-Unaware to AI-Native. Examinable across Modules 1, 5, and 7.*

1.2 Organizational Design for AI Integration

Learning Outcome Statements (LOS)

- Describe how AI adoption changes team structures, role definitions, and reporting hierarchies in investment firms.

- Explain the evolving role of the portfolio manager, analyst, and risk officer in AI-augmented workflows.
- Design governance structures appropriate for an AI-integrated investment firm.

Required Knowledge

- Emergence of the AI-augmented investment team: AI Oversight Lead, Prompt Engineer, Data Architect (Financial Signals)
- The portfolio manager's role: contracting dimensions (data processing, pattern recognition, screening) against expanding dimensions (contextual interpretation, fiduciary accountability, novel situation navigation, AI governance)
- Fiduciary Intelligence: the human capacity to apply ethical reasoning and professional responsibility to AI-assisted decisions
- AI governance structures: AI Investment Policy Statement, Model Risk Management framework, AI Oversight Committee, model validation cycles
- The 'centaur model': designing optimal human-AI teaming arrangements for specific task types
- Communicating AI-augmented investment processes to clients, consultants, and regulators

1.3 Human-AI Teaming in Investment Management

Learning Outcome Statements (LOS)

- Distinguish between the four human-AI teaming patterns and evaluate which is appropriate for a given investment task.
- Identify cognitive risks in human-AI teams and propose governance responses.
- Describe practices that maintain independent human judgment in AI-augmented investment processes.

Required Knowledge

- **The Human-AI Teaming Spectrum, four patterns:**
  - AI as Research Assistant: human directs, AI executes bounded tasks; appropriate for high-stakes or variable-reliability contexts
  - AI as Analytical Co-Pilot: iterative human-AI dialogue; effective for complex research requiring breadth and depth
  - AI as Process Orchestrator: AI manages a multi-step workflow with human review at checkpoints; common in systematic strategies
  - AI as Autonomous Agent: AI operating within defined parameters; appropriate only for execution tasks and never for investment decisions
- **Cognitive risks:**
  - Automation Bias: over-reliance on AI outputs, most dangerous when the AI is performing well because vigilance declines
  - Algorithm Aversion: systematic under-reliance on AI following salient failure events
  - Capability Illusion: organizational belief in AI capability without genuine structural integration
- Practices for maintaining judgment: structured pre-work, mandatory challenge documentation, rotation through non-AI workflows, red team exercises

1.4 The Investment Process in the Age of AI

Learning Outcome Statements (LOS)

- Describe how AI transforms each phase of the investment process: research, portfolio construction, execution, and reporting.
- Apply the CFIA AI-Augmented Investment Loop to analyze how AI-native firms build compounding informational advantages.

Required Knowledge

> **CFIA PROPRIETARY FRAMEWORK™ | CFIA AI-Augmented Investment Loop™**
>
> *Six-stage continuous investment cycle (Data → Hypothesis → Simulation → Decision → Feedback → Update). The feedback mechanism is what distinguishes genuinely AI-integrated investment processes from tool-based approaches. Examinable in Modules 1, 3, and 4.*

- **The six stages:**
  - Stage 01 Data: ingestion, validation, enrichment; data quality gates prevent corrupt signals propagating
  - Stage 02 Hypothesis: AI research agents surface candidate theses, which human analysts challenge and prioritize
  - Stage 03 Simulation: AI-assisted backtesting, scenario analysis, portfolio impact modeling
  - Stage 04 Decision: the irreducible human checkpoint, since fiduciary accountability cannot be delegated to AI
  - Stage 05 Feedback: systematic outcome capture, which drives the compounding advantage of the Loop
  - Stage 06 Update: model refinement, signal weight adjustment, knowledge system enrichment
- Research and idea generation: coverage expansion, unstructured data processing, cross-asset correlation, real-time synthesis
- Portfolio construction: factor exposure optimization, constraint management, and scenario analysis, subject throughout to an explainability requirement
- Execution: adaptive algorithms, transaction cost optimization, post-trade feedback
- Reporting and client communication: AI-generated content must be reviewed by a named accountable professional

> **CFIA PROPRIETARY FRAMEWORK™ | CFIA AI Investment Stack™**
>
> *Five-layer architecture running from Layer 1 (Data Infrastructure) to Layer 5 (Portfolio Decision Engine). Each layer depends on the integrity of the layer below. Proprietary advantage originates at Layers 1 and 2. Examinable across Modules 1, 3, 4, and 6.*

MODULE 2 | Exam Weight: 13% | Estimated Study Time: 7–9 hours

Foundation Models & AI Systems

This module addresses the technical landscape of modern AI at the strategic level. Candidates do not need to code, but must be able to make informed decisions about AI systems, including selecting vendors, designing architectures, and evaluating capabilities and failure modes.

2.1 How Modern AI Systems Work

Learning Outcome Statements (LOS)

- Describe the conceptual architecture of large language models, multimodal models, and reasoning agents at a level sufficient for strategic decision-making.
- Explain the key technical concepts underlying modern AI: tokens, context windows, embeddings, attention, temperature, fine-tuning, and retrieval-augmented generation.
- Trace the evolution of AI from rule-based systems through statistical machine learning to foundation models, and articulate why foundation models represent a qualitative shift.

Required Knowledge

- AI evolution timeline: rule-based systems, statistical machine learning, deep learning, transformer architecture (2017), foundation models
- Transformer architecture: attention mechanism, parallel processing, emergent capabilities at scale
- Foundation models: trained once at scale and adapted cheaply, producing a qualitative shift in economics and generality
- Tokens, context windows, and their practical implications for investment applications
- Embeddings: dense vector representations enabling semantic search and RAG systems
- Temperature: controls output predictability, with lower settings appropriate for professional investment applications
- Absence of memory between sessions, stochastic outputs, and training data cutoff: operational implications
- Hallucination: confident generation of factually incorrect content as an inherent property of generative architectures

2.2 AI Design Patterns for Investment Applications

Learning Outcome Statements (LOS)

- Distinguish between the major AI design patterns (prompt engineering, RAG, fine-tuning, and agents) and evaluate which is appropriate for a given investment application.
- Describe the architecture and operational logic of AI agent frameworks, including tool use, memory systems, and multi-agent orchestration.

Required Knowledge

- Prompt engineering: context injection, instruction design, few-shot examples; the lowest barrier to adoption and the highest interpretability risk
- Retrieval-Augmented Generation (RAG): vector databases, embedding-based retrieval, context injection; addresses the training cutoff limitation
- Fine-tuning: domain adaptation on proprietary data, with higher cost, improved consistency, and additional governance requirements
- AI agents: tool use, memory systems, multi-step reasoning, multi-agent orchestration
- Selecting the appropriate design pattern: task complexity, reliability requirements, governance constraints
- Agentic systems in investment: research automation, portfolio monitoring, compliance surveillance

2.3 AI Failure Modes in Investment Contexts

Learning Outcome Statements (LOS)

- Identify the primary failure modes of AI systems in investment contexts and evaluate their risk implications.

Required Knowledge

- Hallucination: generation of false information with high apparent confidence, together with detection and mitigation strategies
- Distributional shift: a model trained on past data encounters a regime not represented in training, and performance degrades
- Prompt injection: adversarial inputs that override system instructions, with security and governance implications
- Context contamination: irrelevant or misleading context degrades output quality
- Overconfidence: calibration failure, in which the model expresses high certainty on low-certainty outputs
- Failure mode interaction: how multiple failure modes compound in production investment systems

2.4 AI Technology Landscape & Vendor Ecosystem

Learning Outcome Statements (LOS)

- Evaluate the AI vendor landscape relevant to investment management, and compare foundation model providers with AI platform vendors.
- Estimate total cost of ownership for AI systems across compute, licensing, integration, maintenance, and governance dimensions.

Required Knowledge

- Frontier model providers: capabilities, pricing, use-case fit, and the strategic tradeoffs between closed-source and open-source models
- AI infrastructure layer: compute providers, vector databases, orchestration platforms, observability tools
- Investment-specific AI platforms: research co-pilots, portfolio analytics, compliance automation, execution intelligence
- Total cost of ownership: compute, licensing, integration engineering, ongoing maintenance, governance overhead
- Vendor concentration risk: dependency on frontier providers, and open-source optionality as a mitigant
- Evaluating vendor claims: benchmark limitations, and task-specific against general performance

MODULE 3 | Exam Weight: 16% | Estimated Study Time: 8–10 hours

Investment Data Infrastructure

Data architecture is the foundation on which all AI-driven investment capability is built, and it is the most common point of failure in investment organizations that attempt AI at scale. Data infrastructure occupies Layer 1 of the CFIA AI Investment Stack™, so the integrity of every layer above it depends on the controls applied here.

3.1 Market Data Infrastructure

Learning Outcome Statements (LOS)

- Describe the components of a professional market data infrastructure and evaluate fitness for AI use cases.
- Distinguish between data types relevant to AI investment applications and assess their governance requirements.

Required Knowledge

- Market data taxonomy: price and volume data, fundamental data, alternative data, derived data
- Real-time and batch data architectures: event streaming, Lambda and Kappa architectures, latency requirements
- Time-series databases and storage architecture: tick data, OHLCV, corporate actions adjustment
- Data normalization and cross-asset harmonization, and the hidden complexity of multi-asset AI systems
- API architecture for market data: REST, WebSocket, and FIX protocol considerations
- Data vendor management: SLA requirements, failover, reconciliation

3.2 Research Data Platforms

Learning Outcome Statements (LOS)

- Design a research data infrastructure appropriate for AI-driven investment research.
- Evaluate alternative data sources for AI use and assess their fitness, bias risks, and governance requirements.

Required Knowledge

- Document ingestion systems: earnings transcripts, regulatory filings, news, patent applications
- Unstructured data processing pipelines: OCR, parsing, and chunking strategies for LLM consumption
- Feature stores: centralized repositories for computed financial features, preventing leakage and duplication
- Reproducibility of research results: version control for datasets, code, and model artifacts
- Experiment tracking: logging research hypotheses, parameters, and outcomes for auditability
- Alternative data integration: ingestion pipelines, vendor onboarding, fitness assessment for AI use
- Satellite imagery, web traffic, ESG data, and geolocation: AI processing requirements and bias risks

3.3 Knowledge Systems for Investment Intelligence

Learning Outcome Statements (LOS)

- Evaluate systems that organize financial knowledge for AI-driven analysis, and design a knowledge infrastructure appropriate for a research-intensive investment firm.

Required Knowledge

- Vector databases: embedding storage, similarity search, and hybrid retrieval, in managed and self-hosted deployments
- RAG architecture for investment research: document chunking, embedding strategies, retrieval optimization
- Internal research repositories: structured storage of proprietary analysis, model outputs, decision rationale
- Domain-specific ontologies: financial entity recognition, relationship mapping, knowledge graphs
- Knowledge system governance: access controls, version management, deprecation policies

3.4 Data Governance and Provenance

Learning Outcome Statements (LOS)

- Apply data governance principles to AI investment systems and evaluate a data governance framework for adequacy.
- Apply the CFIA Data Provenance Chain to assess the integrity of AI-generated investment signals.

Required Knowledge

> **CFIA PROPRIETARY FRAMEWORK™ | CFIA Data Provenance Chain™**
>
> *Six-stage governance standard: Origination → Validation → Transformation → Storage → Retrieval → Deprecation. Each stage requires documentation and accountability assignment. Examinable in Modules 3 and 5.*

- Data lineage: tracking the complete history of a data point from source to model input
- Data quality dimensions: accuracy, completeness, consistency, timeliness, uniqueness, validity
- Metadata management: documentation standards for AI-consumed data assets
- Data access controls and entitlement management for AI systems
- Regulatory data retention requirements for AI-generated investment decisions
- Third-party data risk: vendor dependency, data quality SLAs, and the regulatory compliance of data sources

3.5 Data Quality and Signal Integrity

Learning Outcome Statements (LOS)

- Identify data quality failure modes that corrupt AI model outputs, and design monitoring systems to detect and remediate them.

Required Knowledge

- Survivorship bias: exclusion of failed entities distorts historical datasets and backtest results

- Look-ahead bias: inadvertent use of future information in historical model training or testing
- Dataset contamination: training data pollution corrupts model outputs in production
- Feature leakage: features that carry information not available at prediction time
- Dataset drift: statistical properties change over time, degrading model performance
- Data quality scorecards: frameworks for monitoring and reporting on data asset fitness
- Point-in-time data: requirements for bias-free historical analysis, and the role of vendor point-in-time databases

MODULE 4 | Exam Weight: 18% | Estimated Study Time: 9–11 hours

AI for Alpha Generation & Research

This module addresses the most commercially significant application of AI in investment management, which is the pursuit of excess returns. Candidates must understand how AI systems generate, test, and operationalize investment insights, and how to distinguish genuine signal from spurious pattern-matching.

4.1 AI-Assisted Fundamental Research

Learning Outcome Statements (LOS)

- Evaluate how AI augments fundamental investment research across the research workflow.
- Design an AI-assisted research process that maintains human judgment at critical decision points.

Required Knowledge

- NLP for financial text: earnings call analysis, regulatory filing extraction, news synthesis, sentiment scoring
- AI-assisted financial modeling: automated ratio analysis, peer comparison, anomaly detection in financials
- Analyst workflow redesign: AI as research assistant against AI as analytical co-pilot in fundamental contexts
- The information interpretation advantage: AI expands breadth while human judgment provides depth and context
- Alternative data integration into fundamental research: satellite, web traffic, ESG controversy
- Validation requirements for AI-generated fundamental insights: source verification and hallucination detection

4.2 Quantitative Signal Generation and Factor Research

Learning Outcome Statements (LOS)

- Evaluate machine learning methods for factor research and assess their applicability to investment signal generation.
- Apply the CFIA Signal Integrity Framework to validate AI-generated signals before deployment.

Required Knowledge

> **CFIA PROPRIETARY FRAMEWORK™ | CFIA Signal Integrity Framework (SIF)™**
>
> *Five mandatory deployment tests for AI-generated investment signals: Statistical Validation, Walk-Forward Testing, Regime Analysis, Capacity Analysis, and Live Pilot. A signal that fails any test must not be deployed. Examinable in Modules 4 and 5.*

- Machine learning methods in factor research: gradient boosting, neural networks, Gaussian processes
- Non-linear factor models: capturing regime-dependent relationships and interaction effects

- Signal validation: the distinction between in-sample and out-of-sample performance
- Walk-forward testing and out-of-sample analysis: methodology and common failure patterns
- Multiple testing problem: how factor research generates false discoveries through excessive specification search
- Regime analysis: validating signal performance across different market environments (bull, bear, crisis, recovery)
- Capacity analysis: position-sizing constraints, market impact, and signal degradation at scale
- Live pilot requirements: minimum pilot duration, performance thresholds, shutdown criteria
- Factor decay and signal alpha half-life: ongoing monitoring requirements after deployment

4.3 AI in Portfolio Construction

Learning Outcome Statements (LOS)

- Evaluate AI applications in portfolio construction and assess their governance implications.

Required Knowledge

- AI-assisted portfolio optimization: Black-Litterman extensions, robust optimization, machine learning methods
- Constraint handling: transaction costs, liquidity, tracking error, ESG mandates, regulatory limits
- Scenario generation with generative AI: tail risk coverage and stress scenario design
- Interpretability in portfolio construction: the obligation of the CIO to explain AI-driven allocations
- Dynamic rebalancing systems: the governance distinctions between automated, AI-recommended, and AI-assisted rebalancing
- ESG integration through AI: scoring systems, controversy monitoring, regulatory compliance tracking

4.4 AI in Trading and Execution

Learning Outcome Statements (LOS)

- Describe AI applications in algorithmic trading and execution, and evaluate their governance and regulatory implications.

Required Knowledge

- Algorithmic execution: VWAP, TWAP, and implementation shortfall, together with their AI extensions and adaptive strategies
- AI-driven pre-trade analytics: market impact prediction, optimal execution timing, venue selection
- Transaction cost analysis (TCA) with AI: decomposition, attribution, and strategy optimization
- Post-trade feedback loops: how execution outcomes improve future AI model performance
- Regulatory considerations for AI-driven execution: MiFID II and SEC Rule 15c3-5 (Market Access Rule)
- Kill switches and circuit breakers: mandatory controls for AI execution systems

MODULE 5 | Exam Weight: 20% | Estimated Study Time: 10–12 hours

Risk, Compliance & AI Governance

This is the highest-weighted module in the CFIA examination, reflecting the centrality of governance to responsible AI deployment. Candidates must be able to design, evaluate, and critique AI governance frameworks at a level well beyond description.

5.1 Model Risk Management in the AI Era

Learning Outcome Statements (LOS)

- Apply model risk management principles to AI systems, and evaluate the limitations of traditional MRM frameworks when applied to foundation models.
- Design a tiered model risk framework that addresses both quantitative models and AI language systems.

Required Knowledge

- SR 11-7 (Federal Reserve): foundational MRM guidance covering model definition, development, validation, and ongoing monitoring
- Why SR 11-7 is insufficient for foundation models: stochastic outputs, black-box architecture, hallucination, and the absence of a deterministic validation path
- AI model risk tiering: risk classification by use case, autonomy level, and consequence of failure
- Model validation for AI systems: adversarial testing, red teaming, benchmark evaluation, sensitivity analysis
- Model documentation standards: model cards, system cards, acceptable use policies
- Ongoing model monitoring: drift detection, performance degradation, distributional shift
- Model inventory and governance registry: tracking AI systems across the investment enterprise

5.2 Explainability, Transparency, and Audit

Learning Outcome Statements (LOS)

- Evaluate explainability techniques for AI investment systems and assess their adequacy for regulatory and fiduciary purposes.

Required Knowledge

- SHAP (Shapley Additive Explanations): feature importance attribution for machine learning models
- LIME (Local Interpretable Model-agnostic Explanations): local approximations for complex models
- Attention visualization and saliency maps for transformer-based models
- The tradeoff between explainability and performance, and when simpler models are governance-preferable
- Audit trail requirements: what must be logged for an AI-assisted investment decision to be defensible

- The inadequacy of output-only logging, and why input, model version, and human interaction must also be captured
- Client disclosure obligations for AI-assisted advice and portfolio management

5.3 The Regulatory Landscape for AI in Financial Services

Learning Outcome Statements (LOS)

- Describe the regulatory framework for AI in financial services across major jurisdictions.
- Evaluate the compliance implications of specific AI deployments under relevant regulatory frameworks.

Required Knowledge

- EU AI Act: risk classification system, prohibited applications, and high-risk category requirements in financial services
- SEC guidance on AI in investment advisory: marketing rule, Investment Advisers Act implications, disclosures
- MiFID II: algorithmic trading requirements, best execution obligations, audit trail standards
- Financial stability concerns: systemic risk from AI model homogeneity, correlated failures, flash crash risk
- IOSCO principles for AI in capital markets: international convergence and divergence
- FSB guidance on AI and financial stability: macroprudential dimensions
- RegTech and AI for compliance: surveillance, monitoring, and reporting automation, together with their own governance requirements

5.4 AI Governance Frameworks

Learning Outcome Statements (LOS)

- Design a comprehensive AI governance framework for an investment management firm.
- Apply the CFIA AI Governance Trident to evaluate an existing governance structure.

Required Knowledge

> **CFIA PROPRIETARY FRAMEWORK™ | CFIA AI Governance Trident™**
>
> *Three-pillar governance structure: Model Risk Management (pillar 1), Explainability and Audit (pillar 2), Regulatory Alignment (pillar 3). An adequate governance framework must address all three pillars simultaneously. Examinable in Module 5.*

- AI Investment Policy Statement: principles for AI use, acceptable use cases, accountability assignments
- AI Oversight Committee: composition, authority, escalation procedures, reporting lines
- The 10 Governing Principles of the CFIA: the ethical and operational framework for AI in investment management
- Human accountability is non-negotiable, because AI systems cannot bear fiduciary responsibility and accountability always rests with a named human
- Incident management for AI failures: detection, escalation, root cause analysis, remediation, disclosure
- Near-zero override rate as a governance red flag: distinguishing genuine human oversight from rubber-stamping

- AI governance documentation: policies, procedures, version control, regulatory examination readiness

5.5 Cybersecurity and Data Privacy in AI Systems

Learning Outcome Statements (LOS)

- Identify cybersecurity risks specific to AI investment systems and evaluate appropriate mitigations.

Required Knowledge

- Model theft and adversarial attacks: extraction attacks, adversarial examples, data poisoning
- Prompt injection in production AI systems: attack vectors and defense architectures
- Privacy risks in AI training data: PII exposure, GDPR implications, data minimization principles
- Vendor security due diligence for AI systems: penetration testing, SOC 2 requirements, incident response SLAs
- Data sovereignty and cross-border AI deployment: jurisdictional constraints on model training and inference

MODULE 6 | Exam Weight: 11% | Estimated Study Time: 5–6 hours

Vendor Strategy & Build/Buy Architecture

Investment firms face critical decisions about whether to build AI capabilities internally, purchase solutions from vendors, or combine both approaches. This module equips candidates to navigate these decisions with discipline and strategic clarity.

6.1 The Build / Buy / Partner Decision Framework

Learning Outcome Statements (LOS)

- Apply a structured framework to evaluate build, buy, and partner decisions for AI capabilities in investment management.

Required Knowledge

- Strategic differentiation test: whether a capability confers proprietary advantage or is a commodity
- Build case: capabilities at Layers 1 and 2 of the CFIA AI Investment Stack (data and knowledge systems), where proprietary advantage originates
- Buy case: commodity AI infrastructure, vendor-hosted models, standardized analytical tools
- Partner case: co-development, licensing, and revenue share, appropriate for specialized domain capabilities
- Partnership structures: joint ventures, white-labeling, API integration agreements, data sharing
- Total cost of build: engineering talent, time-to-market, maintenance burden, opportunity cost
- Vendor dependency risk: lock-in, pricing power, capability roadmap misalignment

6.2 AI Vendor Evaluation and Due Diligence

Learning Outcome Statements (LOS)

- Design a vendor evaluation process for AI systems, and assess vendor claims with appropriate skepticism.

Required Knowledge

- Vendor evaluation criteria: technical capability, performance evidence, security posture, regulatory compliance, financial stability
- Benchmark limitations: why published benchmark scores are insufficient for vendor selection
- Task-specific pilot performance as the required minimum for any production deployment decision
- RFP design for AI procurement: functional requirements, governance requirements, SLA specifications
- Reference checking for AI vendors: questions that reveal production performance as distinct from demonstration performance
- Contract provisions for AI vendors: performance SLAs, data rights, audit rights, exit provisions

6.3 Integration Architecture for Investment Platforms

Learning Outcome Statements (LOS)

- Evaluate integration architectures for embedding AI into investment technology ecosystems.

Required Knowledge

- The investment technology stack: OMS, PMS, risk systems, data warehouses, and reporting platforms as integration touchpoints for AI
- API-first integration: REST, GraphQL, and streaming architectures for AI-connected investment systems
- Data pipeline design for AI consumption: ingestion, transformation, feature engineering, delivery
- Legacy system constraints: integrating modern AI into established enterprise environments
- Testing and validation environments: staging, UAT, and production deployment frameworks for AI systems
- Monitoring and observability: performance metrics, drift detection, alerting for production AI systems
- Change management for AI deployment: stakeholder communication, training, adoption measurement

MODULE 7 | Exam Weight: 10% | Estimated Study Time: 5–6 hours

AI Leadership, Ethics & Fiduciary Duty

The final module addresses the responsibilities that accompany the strategic use of AI in investment management. Candidates must demonstrate that they can act as effective AI architects and as responsible stewards of the trust clients place in investment professionals.

7.1 Building the AI Business Case

Learning Outcome Statements (LOS)

- Construct a credible business case for an AI investment initiative, and evaluate AI initiatives using appropriate financial and strategic metrics.

Required Knowledge

- Value drivers for AI in investment management: research efficiency, signal quality, risk coverage, execution performance, operational scalability
- AI initiative categorization: quick wins, strategic bets, foundational investments, exploratory pilots
- Financial modeling for AI business cases: NPV, IRR, and payback period, together with their limitations for AI initiatives
- Strategic value metrics: competitive positioning, talent acquisition, regulatory positioning, optionality
- AI initiative prioritization: scoring frameworks that combine strategic value, feasibility, and risk
- Board and LP communication: presenting AI strategy with appropriate confidence and transparency

7.2 AI Strategy for Investment Leadership

Learning Outcome Statements (LOS)

- Design a multi-year AI strategy for an investment management firm, and evaluate strategic positioning options.

Required Knowledge

- AI strategy positioning: the tradeoffs between fast follower and first mover approaches in investment AI
- Talent strategy for AI-integrated investment firms: hiring, training, retention, culture
- Multi-year AI roadmap design: sequencing initiatives from foundational infrastructure to competitive differentiation
- Change management and culture: overcoming resistance, building AI literacy, measuring adoption
- Communicating AI strategy to LPs and clients: disclosure obligations and competitive sensitivity
- Board-level AI governance: oversight responsibilities, reporting requirements, education obligations

7.3 The Ethics of AI in Investment Management

Learning Outcome Statements (LOS)

- Apply core AI ethics frameworks to investment management contexts, and evaluate the ethical dimensions of specific AI deployments.

Required Knowledge

- FATE framework: Fairness, Accountability, Transparency, and Explainability applied to investment AI
- Bias in AI investment systems: sources (training data, feature selection, objective function), consequences, mitigation
- Algorithmic discrimination: fair lending, credit allocation, and investment access concerns
- The ethics of AI-driven information asymmetry: when AI advantages conflict with market integrity principles
- Environmental responsibility: the carbon and energy footprint of large-scale AI deployment
- Responsible AI principles for investment management: transparency, human oversight, non-maleficence

7.4 Fiduciary Duty in the Age of AI

Learning Outcome Statements (LOS)

- Evaluate the fiduciary implications of AI adoption in investment management, and determine when AI use is consistent with fiduciary obligations.

Required Knowledge

- The fiduciary standard in investment management: duty of loyalty, duty of care, and the duty to act in the client's best interest
- Whether fiduciary duty can be delegated to an AI system, and why the answer under all major legal frameworks is no
- AI-assisted decisions and the 'algorithm made the decision' defense: why it fails, and why reliance on it is a governance red flag
- Documenting AI-assisted decisions for fiduciary compliance: what must be recorded and retained
- Conflicts of interest in AI systems: vendor relationships, fee structures, data sourcing
- The fiduciary obligation to govern AI, and why negligent AI deployment constitutes a breach of duty
- Emerging case law and regulatory expectations on fiduciary duty and AI

Guidance for Candidates

Recommended Knowledge Base

Candidates are expected to have a foundation in investment management principles equivalent to several years of professional experience or completion of a rigorous investment credential. The following knowledge base is assumed:

- Fundamental investment concepts: asset classes, portfolio theory, risk and return
- Financial markets and instruments: equities, fixed income, derivatives, alternatives
- Basic quantitative literacy: statistics, probability, and data interpretation
- Investment management practice: portfolio construction, risk management, client relationships
- Regulatory awareness: understanding of the investment advisory and asset management regulatory environment

The CFIA Official Curriculum

The official preparation resource for the CFIA examination is the CFIA Official Curriculum, published by ICFDT. The curriculum is the textbook counterpart to this Body of Knowledge, providing the substantive content, explanatory frameworks, case studies, and practice questions that allow candidates to master the Learning Outcome Statements defined here.

The six CFIA proprietary frameworks (AI Maturity Framework, AI Investment Stack, AI-Augmented Investment Loop, Data Provenance Chain, Signal Integrity Framework, and AI Governance Trident) are developed in full in the Official Curriculum. Candidates should master these frameworks to the level of structured application.

Examination Preparation

The CFIA examination tests at three cognitive levels: descriptive (recall and articulation), applied (framework and concept use in realistic scenarios), and evaluative (assessment, critique, and recommendation). Approximately 30% of questions are descriptive, 50% applied, and 20% evaluative.

Evaluative questions are the most differentiating. They typically present a scenario in which an investment professional has deployed AI in some manner, and ask candidates to assess whether the approach is adequate, identify the primary risk, or recommend an improvement. Candidates who understand principles and can apply them to novel situations will significantly outperform those who have memorized definitions.

Recertification & Continuing Professional Development

The CFIA designation must be renewed every three years through the completion of 20 hours of approved continuing professional development (CPD) per year. Given the pace of change in AI, ICFDT maintains and updates the official curriculum on a continuing basis, and charterholders are expected to keep their knowledge current.

CPD credit may be earned through attendance at ICFDT-approved events, completion of approved online courses, publication of relevant research, and participation in ICFDT working groups and standard-setting activities.