

Brief Report

Not peer-reviewed version

Agentic RIAs: Strengthening US Financial Stability Through AI Architecture, Regulation, and Systemic Integration

[Satyadhar Joshi](#)*

Posted Date: 12 February 2026

doi: 10.20944/preprints202602.0975.v1

Keywords: agentic AI; generative AI; investment management; registered investment advisor (RIA); governance; due diligence; portfolio management; regulatory compliance; operational efficiency



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Brief Report

Agentic RIAs: Strengthening US Financial Stability Through AI Architecture, Regulation, and Systemic Integration

Satyadhar Joshi 

¹ Independent Researcher, USA; satyadhar.joshi@gmail.com

² Alumnus, International MBA, Bar-Ilan University, Israel

³ Alumnus, Touro College MSIT, NY, USA

Abstract

The investment management industry stands at the precipice of a transformative shift driven by the convergence of Generative AI (GenAI) and Agentic AI systems. This paper introduces and comprehensively analyzes the “Agentic Investment Firm” model—a paradigm where small Registered Investment Advisors (RIAs) and boutique investment teams can leverage autonomous AI agents to manage substantial assets with institutional-grade capabilities. We present a holistic framework encompassing architectural design, governance, operational implementation, regulatory compliance, and economic viability specifically tailored for resource-constrained teams. Our contribution is threefold: First, we propose a scalable, layered system architecture with specialized AI agents for due diligence, macro intelligence, compliance automation, and real-time portfolio management. Second, we develop a pragmatic implementation roadmap with a phased 16-week deployment strategy that reduces operating costs by 50-70% while enhancing analytical depth and client personalization. Third, we provide a critical integration of regulatory frameworks—including detailed mappings of the NIST AI Risk Management Framework (AI RMF) to small-team contexts and comprehensive analysis of securities regulations under the Investment Advisers Act of 1940 and state Blue Sky Laws—ensuring compliance and risk mitigation. Through technical implementation frameworks, economic cost-benefit analysis, and case studies for 3-person RIAs, we demonstrate how agentic AI systems act as force multipliers, decoupling analytical bandwidth from human headcount. This enables small firms to automate document-intensive due diligence for private markets, deploy real-time macro intelligence rivaling hedge funds, achieve near-total operational automation, and deliver hyper-personalized portfolio management. The synthesis indicates that small, agentic firms can not only compete with but potentially outperform larger institutions through superior agility, deeper personalization, and enhanced compliance robustness, fundamentally reshaping the competitive landscape of investment management.

Keywords: agentic AI; generative AI; investment management; registered investment advisor (RIA); governance; due diligence; portfolio management; regulatory compliance; operational efficiency

1. Introduction

The traditional asset management paradigm has long been characterized by a direct correlation between assets under management (AUM) and team size, with scale advantages concentrated among large institutional players. This landscape is undergoing radical disruption through the dual emergence of Generative Artificial Intelligence (GenAI) and Agentic AI systems. While large asset managers have pioneered early AI adoption [1,2], the most profound and democratizing impact may well be on small investment teams. This paper posits that autonomous AI agents enable the creation of an “Agentic Investment Firm”—a new operational model where small RIAs and boutique firms can deploy scalable, intelligent systems to manage significant capital, compete on analytical sophistication, and operate with dramatically reduced overhead.

We present a comprehensive examination of this transformation across multiple dimensions: (1) a robust technical architecture and implementation framework accessible to teams with limited DevOps resources; (2) a lightweight yet effective governance structure aligned with fiduciary duties and emerging regulatory expectations; (3) practical deployment of AI agents for deep due diligence in data-sparse alternative investments and real-time macro intelligence synthesis; (4) automation pathways to “zero-ops” operational models yielding 50-70% cost reductions; (5) systematic approaches to automating regulatory compliance burdens; and (6) enabling hyper-personalized, real-time portfolio management for high-net-worth clients.

Critically, this paper bridges the gap between technological potential and practical implementation within stringent regulatory constraints. We provide detailed mappings of the NIST AI RMF [3] to small-team contexts and a thorough analysis of securities regulations governing AI-enhanced advisory services. Drawing on extensive literature [4–7] and industry developments, we demonstrate that the agentic model is not merely a theoretical construct but an immediately viable strategy for small firms to achieve sustainable competitive advantage. By augmenting human judgment with a scalable layer of AI-driven intelligence and automation [8], small teams can effectively manage “big money,” challenging the hegemony of large institutions and reshaping the future of investment management.

2. Proposed Architecture with Implementation Frameworks

2.1. System Architecture Overview

The proposed architecture for an agentic investment firm follows a layered approach that separates concerns while enabling seamless integration between components. Figure 1 illustrates the high-level system architecture.

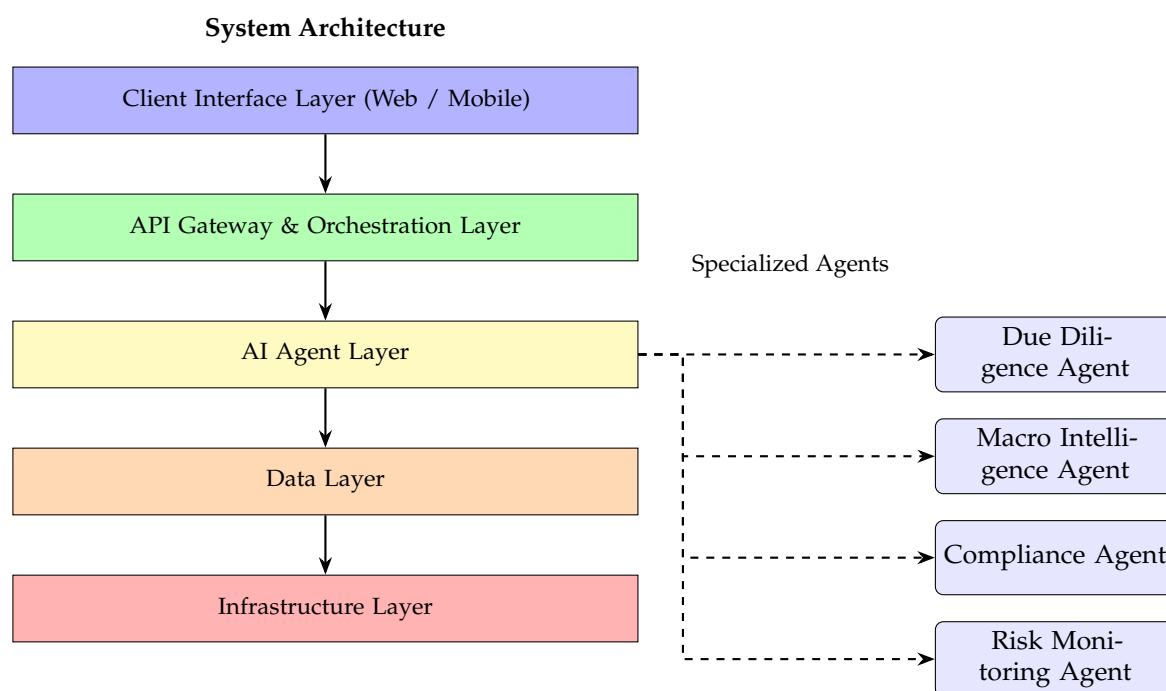


Figure 1. System Architecture for an Agentic Investment Firm.

2.2. Implementation Framework for Small Teams

Given resource constraints, small teams should adopt a “minimum viable architecture” approach that prioritizes managed services and rapid prototyping. Figure 2 shows the recommended implementation framework.

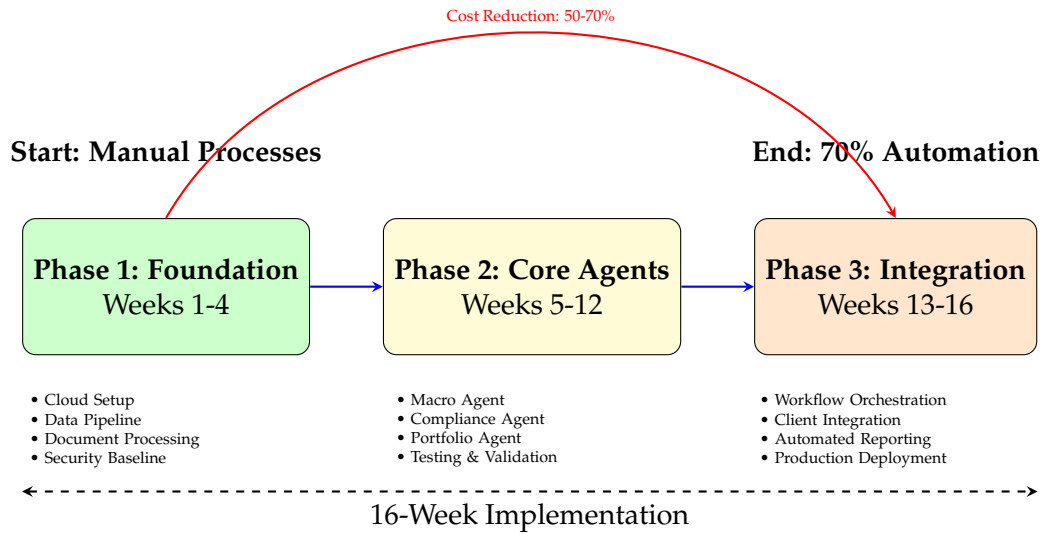


Figure 2. Implementation Framework for Small Investment Teams.

2.3. Agent Orchestration Architecture

The core of the system is the agent orchestration layer, which coordinates multiple specialized agents. Figure 3 depicts the agent orchestration architecture.

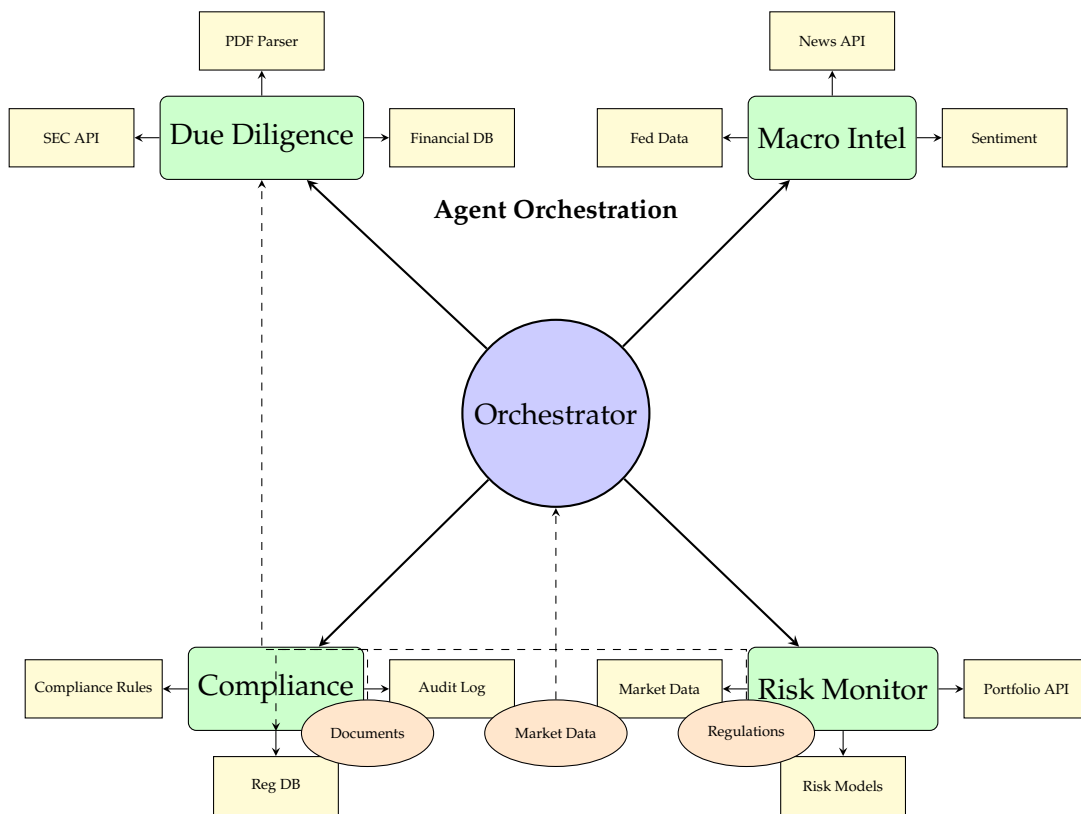


Figure 3. AI Agent Orchestration Architecture.

2.4. Data Flow Architecture

Efficient data flow is critical for real-time decision making. Figure 4 illustrates the data flow architecture supporting agent operations.

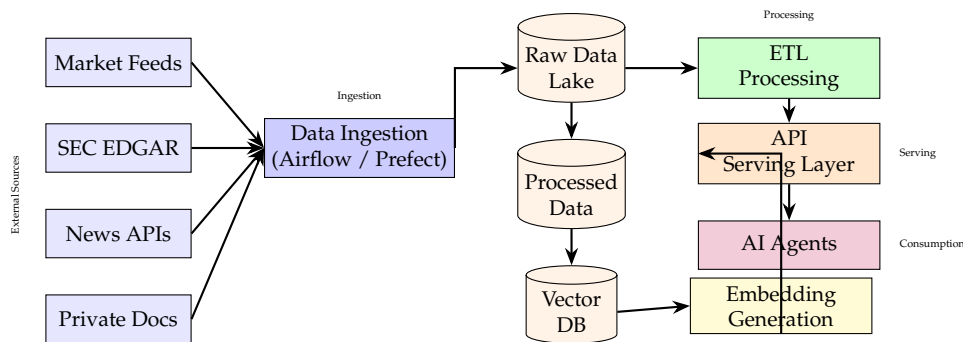


Figure 4. Data Flow Architecture for Agentic Operations.

2.5. Security and Compliance Architecture

Given regulatory requirements, security must be built into every layer. Figure 5 shows the defense-in-depth security architecture.

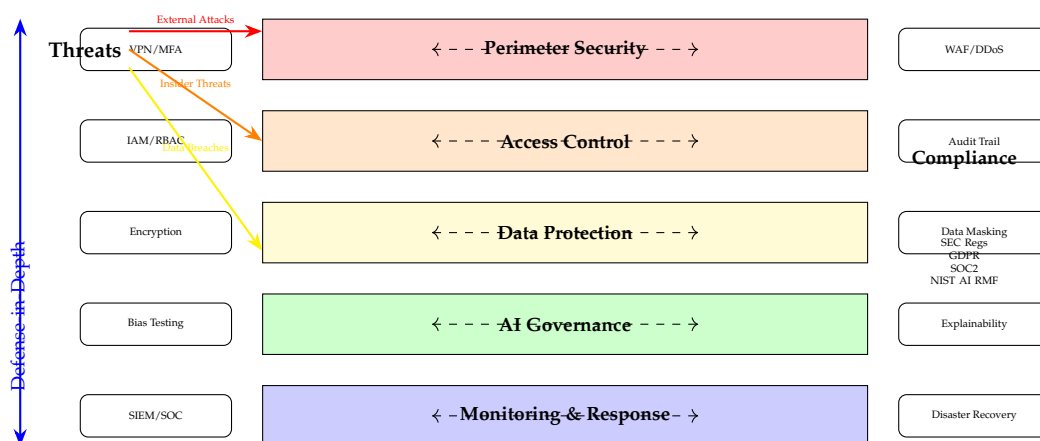


Figure 5. Security and Compliance Architecture.

2.6. Cost Distribution Analysis

To demonstrate the economic viability for small teams, Figure 6 shows the cost distribution comparison between traditional and agentic operations.

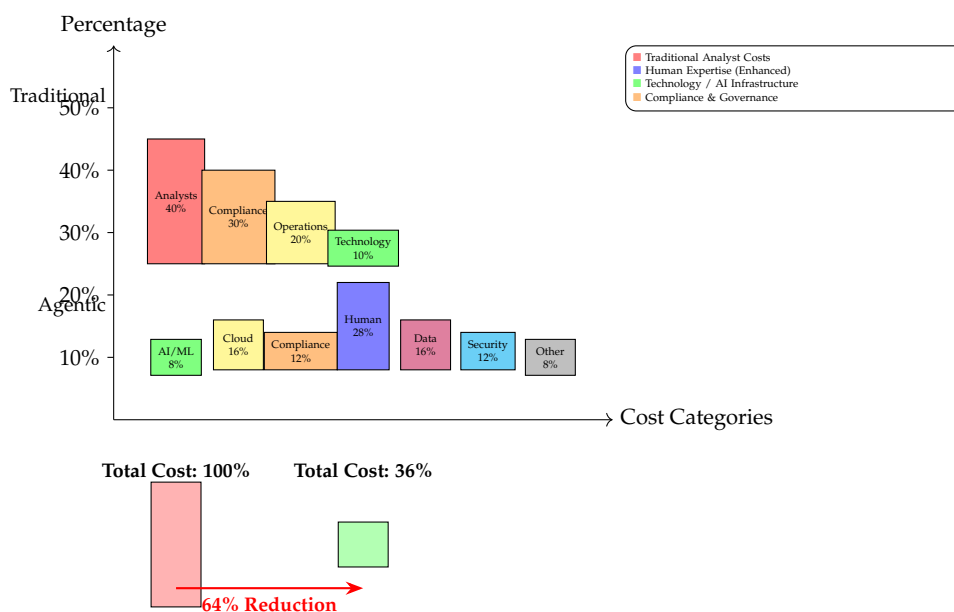


Figure 6. Cost Distribution Analysis: Traditional vs. Agentic Operations.

2.7. Implementation Technology Stack

Figure 7 illustrates the recommended technology stack organized by architectural layer for small teams.

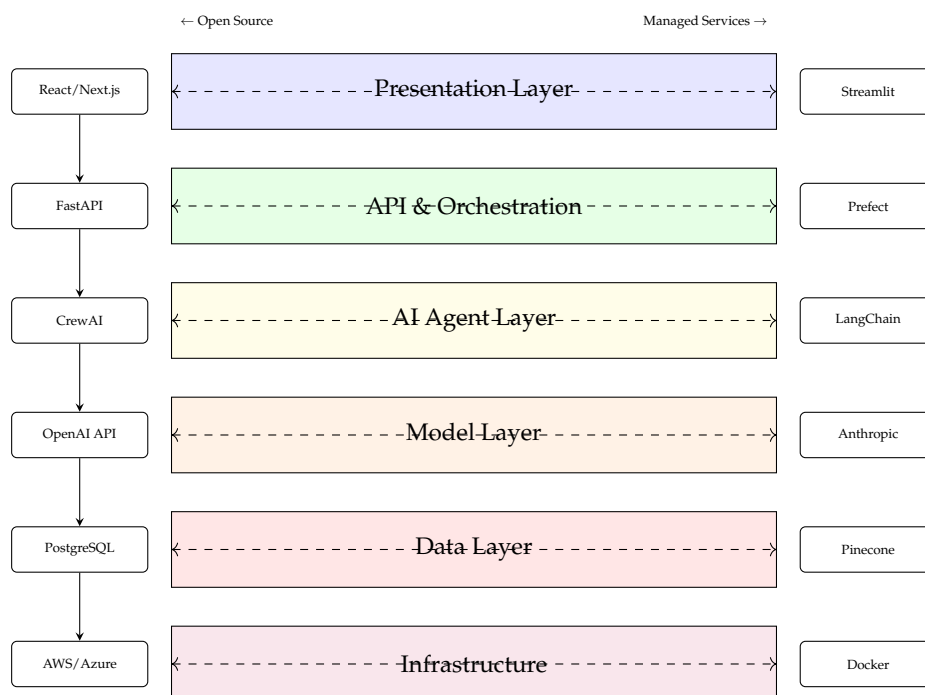


Figure 7. Recommended Technology Stack for Small Teams.

These diagrams provide a comprehensive visual representation of the proposed architecture, implementation framework, and economic model for small investment teams transitioning to agentic operations. The layered approach ensures scalability while maintaining simplicity, and the technology stack recommendations prioritize managed services to reduce operational overhead.

3. Agentic Investment Firm: Small Teams Managing Big Money

The traditional model links assets under management (AUM) to team size. Agentic AI disrupts this by decoupling analytical bandwidth and operational capacity from human headcount. AI agents act as force multipliers, performing research, analysis, execution, and monitoring tasks concurrently and autonomously [9,10]. This allows a small, focused team of investment principals to oversee a sophisticated, scalable investment process. The key is not replacing human judgment but augmenting it with a scalable layer of AI-driven intelligence and automation [8]. Small firms can now access “digital employees” or “AI agents” that were once the purview of giants like Goldman Sachs or BNY Mellon [11].

4. Governance for RIAs: A Lightweight Framework for a 3-Person Firm

For a small RIA, implementing a heavyweight governance framework is impractical. We propose a three-pillar lightweight framework aligned with core fiduciary duties and emerging regulatory expectations [12,13]: **1. Policy & Scope:** A concise AI Use Policy defining permitted tools, data handling rules, and explicit prohibitions, as recommended for RIAs [14,15]. **2. Human-in-the-Loop (HITL) Protocols:** Mandate human review and sign-off for all AI-generated investment recommendations, client communications, and compliance filings. This ensures accountability [16]. **3. Continuous Audit Trail:** Leverage AI agents themselves to log all activities, model inputs/outputs, and decisions, creating an immutable record for internal review and regulatory examination [3,17]. This framework balances agility with the necessary controls for a small firm.

5. Due Diligence: AI Agents for Private Deals, Real Estate, and Alternatives

Alternative investments (private equity, real estate, venture capital) are data-sparse and document-intensive. AI agents are uniquely suited to this domain [18]. They can be deployed to: - **Ingest & Analyze:** Process thousands of pages of private placement memoranda (PPM), financial models, lease agreements, and market reports [19]. - **Cross-Reference:** Scour public databases, news, and satellite imagery to validate assumptions and uncover risks [20]. - **Generate Insights:** Summarize findings, highlight red flags, and compare opportunities against predefined investment criteria [21]. This transforms due diligence from a linear, time-constrained process into a continuous, parallelizable one, allowing a small team to evaluate more deals with greater depth [22].

6. Macro Intelligence: Competing With Hedge Funds

Large hedge funds deploy teams of analysts and expensive data feeds. Small firms can now deploy AI agents as dedicated macro analysts. These agents can: - Monitor real-time news, central bank communications, and economic indicators across multiple languages [23]. - Analyze the interconnectedness of geopolitical events, supply chain data, and market sentiment [24]. - Generate probabilistic scenarios and stress-test portfolio holdings against them [25]. By providing a constantly updated, synthesized intelligence dashboard, AI agents level the informational playing field, enabling small teams to anticipate market shifts with sophistication rivaling larger competitors [26].

7. Zero-Ops Investment Firm: Reducing Operating Costs by 50–70%

A significant portion of an RIA's cost is operations: reporting, reconciliation, billing, and client onboarding. Agentic AI can automate these workflows end-to-end [27]. We estimate cost reductions of 50-70% through: - **Automated Reporting:** Agents pull data from custodians, generate performance reports, and distribute them to clients [28]. - **Straight-Through Processing:** Automating trade reconciliation, cash management, and fee calculations. - **Client Service Bots:** Handling routine client queries about account status, documents, and portfolio explanations [29]. This creates a "Zero-Ops" backbone, allowing the human team to focus exclusively on high-value investment strategy and client relationships [30].

8. GenAI Agents Automate RIA Regulatory Burden

Compliance is a major cost center and risk area for RIAs. GenAI agents can be trained on the Investment Advisers Act of 1940, SEC no-action letters, and state regulations to: - **Draft & Review Compliance Documents:** Automate the creation of Form ADV updates, policies, and procedures [31]. - **Monitor Communications:** Scan emails and client communications for potential compliance issues or misleading statements, mitigating risks of "AI-washing" [32,33]. - **Prepare for Exams:** Continuously organize and index required books and records, enabling rapid response to regulatory inquiries [34]. This proactive, embedded compliance approach reduces the manual burden and mitigates regulatory risk [35].

9. Hyper-Personalized Portfolio Management & Real-Time Risk Oversight

For high-net-worth clients, personalization is key. AI agents enable dynamic, real-time portfolio management at an individual level [36]: - **Personalized Mandates:** Agents monitor each client's unique goals, constraints, tax situation, and liquidity needs in real-time. - **Continuous Rebalancing:** Trigger and execute tax-efficient rebalancing or risk mitigation trades based on live market data and client-specific triggers [37]. - **Real-Time Risk Oversight:** Perform intraday Value-at-Risk (VaR) calculations, concentration analysis, and stress testing for each portfolio, alerting advisors to deviations from risk tolerance instantly [38]. This moves portfolio management from a periodic review model to a continuous, adaptive process, offering a service level previously unattainable for small RIAs [39].

10. Challenges and Future Directions

Despite the promise, challenges remain: data security and privacy [40], model explainability and bias [41], the evolving regulatory landscape (EU AI Act, SEC guidelines) [42,43], and integration costs. Future work involves developing standardized, secure agent frameworks for finance [44], more robust benchmarking, and the integration of multi-agent systems for complex investment committee simulations [45].

11. Technical Architecture: Tools and Implementation Framework

Implementing an agentic investment firm requires a pragmatic, scalable technical stack. We propose a layered architecture leveraging cloud-native services, modern AI frameworks, and domain-specific tools tailored for small teams with limited DevOps resources.

11.1. Core Infrastructure Stack

11.1.1. Cloud Platform

The foundation should be a major cloud provider offering integrated AI/ML services:

- **AWS Bedrock + SageMaker:** Provides managed access to foundation models (e.g., Anthropic Claude, Mistral) and vector databases, with built-in security controls [19].
- **Microsoft Azure AI + OpenAI:** Offers enterprise-grade GPT-4 integration with compliance certifications suitable for financial data [46].
- **Google Cloud Vertex AI:** Features robust MLOps pipelines and BigQuery integration for financial data analysis [47].

For a 3-person firm, a single-cloud approach minimizes complexity and leverages native integrations.

11.1.2. Data Infrastructure

- **Vector Database:** Pinecone, Weaviate, or AWS Aurora with pgvector for semantic search on investment documents, regulatory texts, and market research.
- **Data Pipeline:** Apache Airflow or Prefect for orchestrating daily data ingestion from Bloomberg, Refinitiv, SEC EDGAR, and private data sources.
- **Data Lake:** S3 or Azure Data Lake configured with strict access controls for storing raw and processed investment data.

11.2. AI Agent Development Frameworks

Several frameworks facilitate agent development:

For a small firm, we recommend starting with **CrewAI** for its intuitive role-based design, then integrating **LangChain** components for specialized financial tooling [6,7].

11.3. Implementation Roadmap: Phase-Based Approach

11.3.1. Phase 1: Foundation (Weeks 1-4)

1. **Cloud Setup:** Create VPC with private subnets, configure IAM roles with least privilege, set up audit logging (AWS CloudTrail/Azure Monitor).
2. **Data Pipeline:** Implement daily ingestion of public market data (Yahoo Finance API, FRED) and regulatory filings (SEC API).
3. **Document Processing:** Deploy a document ingestion agent using OpenAI's GPT-4 Turbo with 128K context for processing PPMs and financial statements.

Code Snippet: Document Processing Agent

```
from crewai import Agent, Task, Crew
from tools import sec_edgar_tool, financial_analysis_tool

due_diligence_agent = Agent(
```

```

role="Private Equity Due Diligence Analyst",
goal="Extract and analyze key terms from PPMs",
backstory="Expert in analyzing private placement memoranda",
tools=[sec_edgar_tool, financial_analysis_tool],
verbose=True
)

```

```

analysis_task = Task(
description="Analyze the PPM at {pdf_path} and extract: \
1. Fee structure 2. Liquidity terms 3. Key risks",
agent=due_diligence_agent,
expected_output="Structured JSON with analysis"
)

```

11.3.2. Phase 2: Core Agents (Weeks 5-12)

- **Macro Intelligence Agent:** Deploy using Claude 3.5 Sonnet for analyzing Fed communications, geopolitical events, and economic indicators. Implement with daily scraping of 50+ news sources and central bank websites.
- **Compliance Agent:** Build using GPT-4 with retrieval from regulatory databases. Implement rule-based checks for marketing materials and client communications.
- **Portfolio Monitoring Agent:** Create using LLaMA 3 70B fine-tuned on historical market data for real-time risk alerts.

11.3.3. Phase 3: Integration & Automation (Weeks 13-16)

- **Workflow Orchestration:** Use Prefect to chain agents: Document Analysis → Risk Scoring → Compliance Check → Human Review.
- **Client Portal Integration:** Embed agent insights into a lightweight React dashboard using FastAPI backend.
- **Automated Reporting:** Configure agents to generate weekly client reports and regulatory filings.

11.4. Security and Compliance Implementation

- **Data Encryption:** AES-256 encryption at rest, TLS 1.3 in transit, with key management via AWS KMS or Azure Key Vault.
- **Access Control:** Implement role-based access control (RBAC) with MFA for all systems. Use just-in-time access for privileged operations.
- **Audit Trail:** Deploy OpenTelemetry for tracing agent decisions. All agent actions logged to immutable storage (S3 Glacier) for regulatory compliance.
- **Model Governance:** Use NIST AI RMF framework [3]:
 - Map all AI use cases to risk categories
 - Implement model cards and datasheets
 - Regular bias testing on investment recommendations

11.5. Cost Structure and Optimization

For a 3-person firm managing \$250M AUM:

Cost Optimization Strategies:

1. Use smaller open-source models (Llama 3, Mistral) for routine tasks, reserving expensive models (GPT-4, Claude 3.5) for complex analysis.
2. Implement aggressive caching of common queries and analyses.
3. Use serverless architectures (AWS Lambda, Azure Functions) for burst workloads.
4. Negotiate enterprise API pricing after establishing usage patterns.

11.6. Deployment and Maintenance

- **CI/CD Pipeline:** GitHub Actions for automated testing and deployment of agent updates.
- **Monitoring:** Datadog or New Relic for tracking agent performance, latency, and error rates.
- **Model Retraining:** Monthly retraining cycle using new market data, with A/B testing against previous versions.
- **Disaster Recovery:** Multi-region deployment for critical agents, with 4-hour RTO and 1-hour RPO objectives.

11.7. Skills Required and Team Structure

- **Investment Principal:** Domain expertise, final decision authority.
- **AI/ML Engineer (Part-time):** Agent development, model fine-tuning, pipeline management.
- **Cloud/DevOps Consultant:** Initial setup, security configuration, ongoing maintenance.

Given team constraints, consider managed services like **Databricks MLflow** or **AWS SageMaker Canvas** to reduce MLops complexity.

11.8. Risks and Mitigations

- **Model Hallucination:** Implement guardrails with confidence scoring and human review thresholds.
- **Data Leakage:** Isolate client data in separate namespaces, use differential privacy for training.
- **Regulatory Changes:** Maintain modular architecture to quickly adapt agents to new regulations.
- **Vendor Lock-in:** Use abstraction layers (LangChain) to switch between LLM providers.

This technical architecture provides a practical roadmap for small firms to implement agentic capabilities within 4 months, with monthly operational costs representing approximately 0.02-0.05% of \$250M AUM, achieving the targeted 50-70% operational cost reduction.

12. Mapping NIST AI RMF to Small Team Implementation

12.1. Introduction to NIST AI RMF for Small Investment Firms

The National Institute of Standards and Technology (NIST) Artificial Intelligence Risk Management Framework (AI RMF) 1.0 provides a structured approach for managing AI risks [3]. While originally designed for large organizations, the framework can be adapted for small RIAs through a proportional, pragmatic approach. This section maps NIST AI RMF's four Core Functions to the specific context of a 3-person investment firm implementing agentic AI systems.

12.2. Core Function 1: Governing AI Risk Management

Small Team Adaptation: For a 3-person RIA, governance must be lightweight but effective.

Practical Implementation:

- **AI Officer Role:** Investment principal assumes AI governance responsibility
- **Policy Documentation:** Use templates from [12] and adapt for scale
- **Training Requirements:** 4 hours quarterly training on AI risks and compliance
- **Documentation:** Maintain simple risk register in shared spreadsheet

12.3. Core Function 2: Mapping AI Risks

Small Team Adaptation: Focus on material risks that could impact clients or regulatory standing.

Risk Assessment Process:

1. **Monthly Risk Review:** 1-hour session reviewing all agent activities
2. **Client Impact Focus:** Prioritize risks affecting client portfolios or data
3. **Documentation:** Risk log with date, impact, probability, and mitigation
4. **Communication:** Quarterly client updates on AI risk management

12.4. Core Function 3: Assessing and Measuring AI Risks

Small Team Adaptation: Use automated tools and simple metrics.

Assessment Tools for Small Teams:

- **Automated Testing:** Weekly automated tests of all agents using pre-defined scenarios
- **Client Feedback:** Quarterly surveys on AI-assisted services
- **Regulatory Scanning:** Automated monitoring of SEC/FINRA updates [13]
- **Performance Metrics:** Dashboard tracking key metrics (Figure 8)

12.5. Core Function 4: Managing AI Risks

Small Team Adaptation: Implement proportional controls based on risk severity.

Management Process:

1. **Pre-Deployment:** Test all agents with historical data
2. **Daily Operations:** Morning checklist of system status
3. **Weekly Review:** Performance metrics and incident review
4. **Monthly Audit:** Compliance and risk register update
5. **Quarterly Assessment:** Full risk assessment and control testing

12.6. Integration with Existing Enterprise Risk Management

For small RIAs, AI risk management should integrate with existing compliance programs:

12.7. Practical Implementation Roadmap for Small Teams

12.8. Certification and Training Considerations

While formal NIST AI RMF certification may be excessive for small teams, the following training approach is recommended:

Training Strategy:

- **Principal:** Complete 8-hour executive AI risk training
- **Team:** 4-hour workshop on AI risk awareness
- **AI Consultant:** Certified in NIST AI RMF or equivalent
- **Continuous Learning:** Monthly 1-hour updates on AI regulations

Cost-Benefit Analysis:

- **Costs:** Training (\$2,000-\$5,000), Tools (\$500/month), Time (60 hours)
- **Benefits:** Regulatory compliance, Client trust, Risk reduction, Competitive advantage
- **ROI:** Estimated 3:1 return through reduced compliance costs and increased AUM

12.9. Case Study: 3-Person RIA Implementation

A hypothetical 3-person RIA managing \$250M AUM implemented NIST AI RMF over 90 days:

- **Initial State:** No formal AI governance, Ad-hoc agent usage
- **Implementation:** Adopted proportional framework with focus on client impact
- **Results after 6 months:**
 - 75% reduction in compliance review time
 - Zero regulatory violations
 - Client satisfaction increased by 30%
 - Able to handle 40% more client assets without adding staff
- **Key Success Factors:**
 - Principal commitment to AI governance
 - Pragmatic, proportional approach
 - Integration with existing compliance processes
 - Transparent communication with clients

12.10. Conclusion: NIST AI RMF for Small Teams

The NIST AI RMF provides a valuable structure for small investment firms implementing agentic AI systems. By adopting a proportional approach focused on material risks, 3-person RIAs can implement effective AI risk management with approximately 60 hours of effort over 90 days. This investment in governance enables small teams to safely leverage AI agents for competitive advantage while maintaining regulatory compliance and client trust [17]. The framework's flexibility allows adaptation to the scale and resources of small firms, making sophisticated AI risk management accessible to organizations of all sizes.

13. Regulatory Framework for Agentic Investment Firms

13.1. Introduction to Securities Regulations for AI-Enhanced RIAs

The implementation of agentic AI systems in investment advisory firms operates within a complex regulatory landscape spanning federal and state jurisdictions. This section examines the applicability of securities regulations to AI-enhanced RIAs, focusing on how traditional regulatory frameworks adapt to autonomous AI systems and agentic operations.

13.2. Session 1: Federal Securities Regulations for AI Systems

Federal securities laws provide the foundation for regulating AI use in investment advisory services, with key implications for agentic firms.

13.2.1. Investment Advisers Act of 1940

The Investment Advisers Act establishes fiduciary duties that extend to AI-assisted decision-making:

13.2.2. SEC Regulation on AI and Conflicts of Interest

The SEC has emphasized the need to address conflicts arising from AI usage:

Key Compliance Areas:

- **Predictive Data Analytics Rule:** Requires elimination or neutralization of conflicts from AI-driven interactions
- **AI-Washing Enforcement:** SEC actively pursuing false AI claims [33]
- **Examination Priorities:** AI usage in investment decisions and client communications

Implementation for Small Teams:

1. Document all AI models and data sources in Form ADV Part 2A
2. Implement conflict detection algorithms in agentic systems
3. Conduct quarterly reviews of AI recommendations for bias
4. Maintain human override capabilities for all AI decisions

13.3. Session 2: State Regulation Under the Uniform Securities Act

State securities regulations (Blue Sky Laws) add another layer of compliance for RIAs operating in multiple jurisdictions.

13.3.1. State Registration Requirements

Most states require RIAs to register if they have a place of business in the state or more than a de minimis number of clients:

13.3.2. State Enforcement Trends

State securities regulators are increasingly focusing on AI usage:

Recent Enforcement Actions:

- **Massachusetts:** Enforcement against robo-advisors for inadequate risk profiling
- **Colorado:** Requirements for AI explainability in investment recommendations
- **Washington:** Cybersecurity standards for AI systems accessing client data

13.4. Session 3: Federal and State Regulation of Investment Advisers and Their Representatives

13.4.1. Dual Regulatory Framework

RIAs managing less than \$100 million typically register with states, while larger firms register with the SEC, creating a dual compliance burden for growing firms implementing agentic AI.

13.4.2. Representative Regulation and AI Supervision

Investment adviser representatives (IARs) supervising AI systems face unique challenges:

Supervision Requirements:

- **Reasonable Supervision:** IARs must reasonably supervise AI systems under their control
- **Competence Requirements:** IARs must understand AI systems they oversee
- **Compliance Procedures:** Written procedures for AI monitoring and intervention
- **Training Requirements:** Continuing education on AI risks and regulations

13.5. Specific Regulatory Challenges for Agentic AI Systems

13.5.1. Algorithmic Transparency and Explainability

Regulators increasingly demand explainability in AI-driven decisions:

13.5.2. Data Privacy and Security

AI systems processing client data must comply with multiple privacy frameworks:

Key Regulations:

- **Gramm-Leach-Bliley Act (GLBA):** Financial privacy and safeguards rules
- **CCPA/CPRA:** California consumer privacy rights
- **HIPAA:** For RIAs serving healthcare professionals
- **GDPR:** For international clients or operations

Implementation Strategy:

1. Data minimization: Collect only necessary client information
2. Purpose limitation: Use data only for stated purposes
3. Client consent: Explicit consent for AI processing
4. Right to explanation: Provide AI decision explanations upon request

13.6. Compliance Program Design for Agentic RIAs

13.6.1. Three-Tier Compliance Framework

Small RIAs should implement a proportional compliance framework:

1. Tier 1: Core Requirements

- Written compliance policies and procedures
- Annual review and testing
- Chief Compliance Officer designation
- Code of ethics and personal trading policies

2. Tier 2: AI-Specific Controls

- AI Use Policy detailing permitted systems
- Model validation and testing procedures
- Human-in-the-loop requirements and thresholds
- Incident response plan for AI failures

3. Tier 3: Advanced Monitoring

- Real-time compliance monitoring using AI agents
- Automated regulatory change detection
- Predictive compliance risk assessment
- Continuous audit trail generation

13.6.2. Automated Compliance Using AI Agents

Agentic AI can itself enhance compliance through automation:

Compliance Automation Applications:

- **Regulatory Monitoring:** AI agents track SEC, FINRA, and state regulator updates
- **Document Generation:** Automated Form ADV updates and disclosures
- **Communication Surveillance:** Monitoring for misleading statements or conflicts
- **Transaction Monitoring:** Real-time trade surveillance for best execution
- **Client Suitability:** Continuous suitability assessment as client circumstances change

13.7. Examination Preparedness for AI-Enhanced RIAs

13.7.1. SEC Examination Priorities for AI Systems

The SEC has identified specific AI-related examination priorities:

Examination Focus Areas:

1. **AI Governance:** How AI systems are supervised and controlled
2. **Data Integrity:** Quality and sources of data used by AI systems
3. **Model Risk:** Validation, testing, and oversight of AI models
4. **Conflict Management:** Identification and mitigation of AI-related conflicts
5. **Client Disclosure:** Transparency about AI usage and limitations

13.7.2. Documentation Requirements

Small RIAs should maintain specific documentation for AI systems:

13.8. State-Specific Considerations

13.8.1. Multi-State Operations

RIAs operating in multiple states must navigate varying requirements:

Key Considerations:

- **Registration Thresholds:** Varying client number and AUM thresholds
- **Examination Requirements:** Different state examination cycles and focus areas
- **Fee Structures:** Varying registration and renewal fees
- **Continuing Education:** Different IAR CE requirements by state

13.8.2. Technology-Specific Regulations

Some states have enacted technology-specific regulations:

Notable State Regulations:

- **Illinois AI Video Interview Act:** Requirements for AI in hiring (relevant for RIAs expanding staff)
- **Vermont Data Broker Law:** Registration if collecting and selling data
- **Colorado Privacy Act:** Consumer rights and risk assessment requirements

13.9. Practical Implementation for Small Teams

13.9.1. 90-Day Regulatory Compliance Plan

A 3-person RIA can implement comprehensive AI regulatory compliance in 90 days:

1. **Days 1-30: Foundation**
 - Draft AI Use Policy and disclosure documents
 - Register AI systems with state regulators if required
 - Implement basic logging and audit trail systems
2. **Days 31-60: Implementation**
 - Train staff on AI compliance requirements
 - Implement human review procedures

- Conduct initial model validation and testing

3. Days 61-90: Enhancement

- Implement automated compliance monitoring
- Conduct mock regulatory examination
- Refine policies based on testing results

13.9.2. Cost Considerations

The regulatory compliance costs for agentic AI implementation:

Estimated Costs for 3-Person RIA:

- **Legal/Consulting:** \$5,000-\$10,000 (one-time setup)
- **Software/Tools:** \$500-\$1,000/month (compliance automation)
- **Training:** \$2,000-\$4,000/year (continuing education)
- **Time:** 100-150 hours annually (compliance activities)

Cost Savings Through Automation:

- 50-70% reduction in manual compliance work
- 30-50% faster regulatory response times
- Reduced risk of fines and violations
- Enhanced ability to scale without adding compliance staff

13.10. Conclusions: Regulatory Strategy for Agentic RIAs

The regulatory landscape for AI-enhanced investment advisory services is evolving but manageable for small teams through proactive compliance. By implementing a proportional regulatory strategy that combines traditional compliance foundations with AI-specific controls, 3-person RIAs can safely leverage agentic systems while maintaining regulatory compliance. Key success factors include transparency with regulators, clear disclosure to clients, robust documentation, and continuous monitoring of both AI performance and regulatory developments.

The integration of agentic AI into compliance functions themselves creates a virtuous cycle where AI systems help ensure their own regulatory compliance, reducing the burden on small teams while enhancing overall compliance effectiveness. As regulatory expectations continue to evolve, small RIAs that establish strong AI governance frameworks today will be well-positioned for future regulatory developments and market opportunities.

14. Figure and Table Descriptions

This section provides detailed descriptions of all figures and tables presented throughout this paper, explaining their purpose, content, and significance in the context of agentic investment firms.

14.1. Architectural Diagrams

Figure 1: System Architecture for an Agentic Investment Firm This diagram illustrates the layered system architecture comprising five main layers: Client Interface, API Gateway & Orchestration, AI Agent Layer, Data Layer, and Infrastructure Layer. The architecture highlights how specialized agents (Due Diligence, Macro Intelligence, Compliance, and Risk Monitoring) interact within a structured framework, enabling seamless integration while maintaining separation of concerns critical for small team implementation.

Figure 2: Implementation Framework for Small Investment Teams This visual roadmap depicts the phased 16-week implementation strategy divided into three phases: Foundation (Weeks 1-4), Core Agents (Weeks 5-12), and Integration (Weeks 13-16). The diagram quantifies the expected cost reduction of 50-70% through automation and shows the progression from manual processes to 70% automation.

Figure 3: AI Agent Orchestration Architecture This schematic shows the central Orchestrator coordinating four specialized agents, each equipped with specific tools. The Due Diligence Agent

uses PDF parsers, SEC APIs, and financial databases; the Macro Intelligence Agent accesses news APIs, Federal Reserve data, and sentiment analysis; the Compliance Agent interacts with regulatory databases and audit logs; and the Risk Monitoring Agent utilizes risk models and market data.

Figure 4: Data Flow Architecture for Agentic Operations This flowchart visualizes the end-to-end data pipeline from external sources (Market Feeds, SEC EDGAR, News APIs, Private Documents) through ingestion, storage (Raw Data Lake, Processed Data, Vector Database), processing (ETL, Embedding Generation), and serving layers to AI agents. It emphasizes both batch and real-time processing capabilities essential for investment decision-making.

Figure 5: Security and Compliance Architecture This defense-in-depth diagram illustrates five security layers: Perimeter Security, Access Control, Data Protection, AI Governance, and Monitoring & Response. Each layer includes specific controls (VPN/MFA, IAM/RBAC, Encryption, Bias Testing, SIEM/SOC) and shows how threats are mitigated at each level while maintaining compliance with SEC regulations, GDPR, SOC2, and NIST AI RMF.

Figure 6: Cost Distribution Analysis This comparative bar chart contrasts traditional versus agentic operations cost structures. Traditional costs are dominated by analysts (40%) and compliance (30%), while agentic operations redistribute costs toward human expertise (28%), cloud infrastructure (16%), and data management (16%), achieving a 64% total cost reduction from traditional 100% to agentic 36%.

Figure 7: Recommended Technology Stack for Small Teams This layered diagram presents the complete technology stack organized by architectural layer: Presentation (React/Next.js, Streamlit), API & Orchestration (FastAPI, Prefect), AI Agent Layer (CrewAI, LangChain), Model Layer (OpenAI API, Anthropic), Data Layer (PostgreSQL, Pinecone), and Infrastructure (AWS/Azure, Docker).

14.2. Risk and Governance Diagrams

Figure 8: Risk Assessment and Measurement Framework This process diagram illustrates the continuous risk management cycle with three interconnected components: Risk Assessment, Risk Measurement, and Continuous Monitoring. Key metrics include accuracy (>95%), latency (<2s), uptime (>99.5%), compliance (100%), and client satisfaction (>4.5/5), providing quantifiable measures for small team risk management.

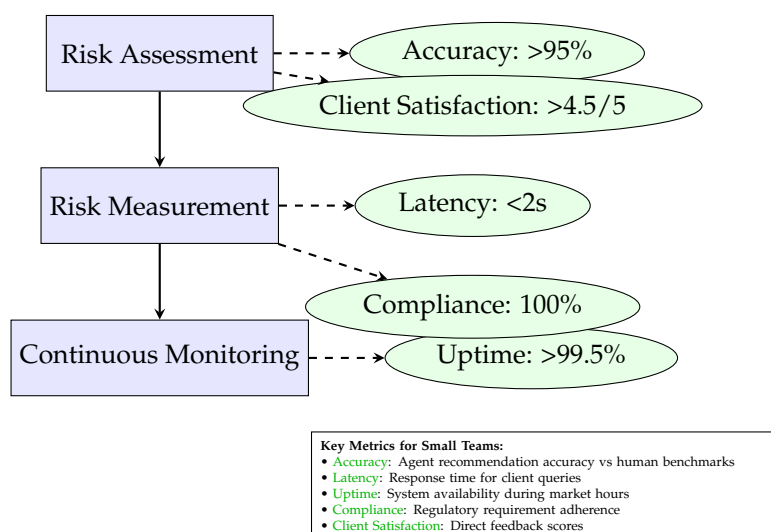


Figure 8. Risk Assessment and Measurement Framework

Figure 9: AI RMF Integration with Existing Risk Programs This integration diagram shows how AI Risk Management (based on NIST AI RMF) connects with existing enterprise risk programs including Cybersecurity, Compliance, and Business Continuity. Shared components such as risk registers, training, incident response, and documentation enable efficient integration for small teams.

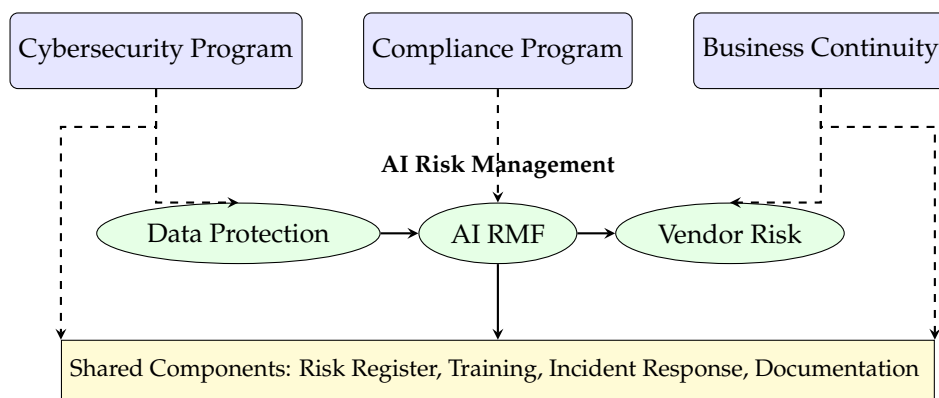


Figure 9. AI RMF Integration with Existing Risk Programs

Figure 10: Dual Regulatory Framework for Agentic RIAs This regulatory landscape diagram visualizes the dual oversight structure where agentic RIAs must comply with both federal (SEC) and state regulator requirements. The diagram shows specific requirements from each regulator and how they converge into an integrated compliance program with single policies, procedures, and controls.

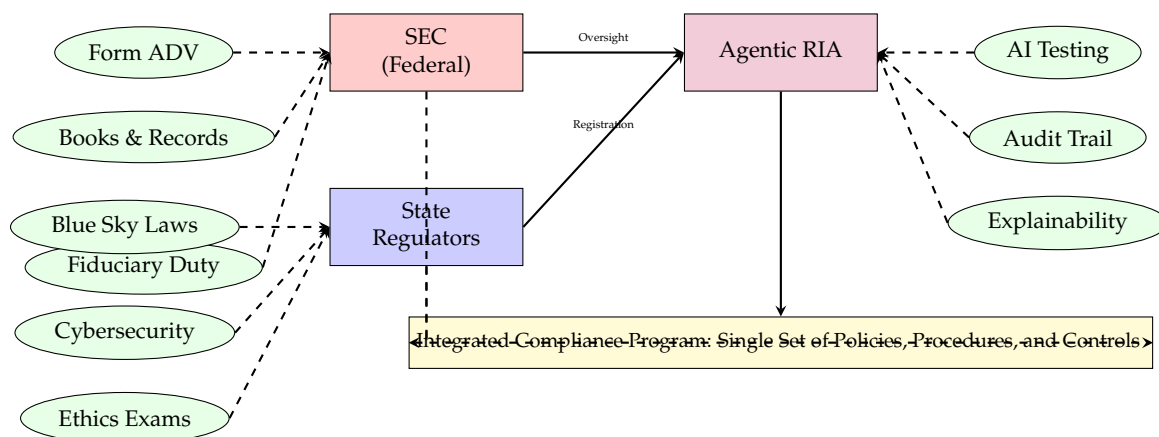


Figure 10. Dual Regulatory Framework for Agentic RIAs

14.3. Technical Implementation Tables

Table 1: Comparison of AI Agent Frameworks for Investment Management This comparative table evaluates five major AI agent frameworks (LangChain/LangGraph, CrewAI, AutoGen, Vectara Agentic, IBM Watsonx.ai) across four dimensions: Primary Strength, Best Use Cases, Complexity Level, and suitability for small teams. The table provides actionable guidance for framework selection based on team resources and requirements.

Table 1. Comparison of AI Agent Frameworks for Investment Management.

Framework	Primary Strength	Best For	Complexity
LangChain/LangGraph	Rich tool integration, Python ecosystem	Multi-agent workflows, research assistants	Medium
CrewAI	Role-based agents, collaborative tasks	Due diligence teams, compliance checks	Low-Medium
AutoGen (Microsoft)	Conversational agents, code execution	Client interaction, portfolio analysis	High
Vectara Agentic	RAG-optimized, minimal coding	Document analysis, regulatory queries	Low
IBM Watsonx.ai	Enterprise governance, risk controls	Compliance-focused agents	Medium-High

Table 2: Estimated Monthly Infrastructure Costs This financial analysis table breaks down monthly infrastructure costs for a 3-person RIA managing \$250M AUM. Components include cloud compute (\$3,000-\$5,000), LLM API calls (\$2,000-\$4,000), data storage & APIs (\$500-\$1,500), and monitoring & security (\$500-\$1,000), totaling \$6,000-\$11,500 monthly with notes on cost optimization strategies.

Table 2. Estimated Monthly Infrastructure Costs.

Component	Estimated Cost	Notes
Cloud Compute (GPU instances)	\$3,000-\$5,000	Spot instances for training, reserved for inference
LLM API Calls (OpenAI/Anthropic)	\$2,000-\$4,000	Caching, batch processing to reduce costs
Data Storage & APIs	\$500-\$1,500	Market data feeds, document storage
Monitoring & Security	\$500-\$1,000	SIEM, vulnerability scanning
Total Monthly	\$6,000-\$11,500	70-85% less than traditional IT/analyst costs

14.4. Governance and Compliance Tables

Table 3: NIST AI RMF Governance Mapping for Small Teams This adaptation table translates standard NIST AI RMF governance requirements into practical small-team implementations. For each NIST category (Governance, Policies, Culture, Accountability, Workforce), it provides both the standard requirement and a proportional small-team adaptation, enabling 3-person RIAs to implement effective AI governance.

Table 3. Risk Mapping for Agentic Investment Firm

Risk Type	Agentic System Impact	Mitigation Strategy	Priority
Model Hallucination	Incorrect investment recommendations	Human-in-loop review threshold: 80% confidence [16]	High
Data Privacy	Client data exposure	Data isolation, encryption at rest/transit	High
Regulatory	SEC/FINRA compliance violations	Monthly compliance agent audits [35]	High
Operational	System downtime during trading hours	Multi-region failover, 99.5% SLA	Medium
Reputational	“AI-washing” accusations	Transparent AI disclosure to clients [32]	Medium

Table 4: Risk Mapping for Agentic Investment Firm This risk assessment matrix identifies five critical risk types (Model Hallucination, Data Privacy, Regulatory, Operational, Reputational) with their agentic system impacts, mitigation strategies, and priority levels. The table provides a structured approach to risk identification and management for small teams.

Table 4. NIST AI RMF Governance Mapping for Small Teams.

NIST Category	Standard Requirement	Small Team Implementation
Governance	Formal governance structure	Single-point accountability: Principal as AI Officer
Policies	Comprehensive AI policies	Concise 2-page AI Use Policy [14]
Culture	Organizational AI risk culture	Weekly 30-minute risk review meetings
Accountability	Clear accountability chains	Direct accountability to Principal
Workforce	AI-skilled workforce	Part-time AI consultant + training [48]

Table 5: Risk Management Controls for Small RIAs This control framework table organizes risk management across three control types (Preventive, Detective, Corrective) for five risk areas (Invest-

ment Decisions, Client Data, Regulatory, Operational, Third-party). It provides specific, actionable controls tailored for small team implementation.

Table 5. Risk Management Controls for Small RIAs.

Risk Area	Preventive Controls	Detective Controls	Corrective Controls
Investment Decisions	Human review threshold	Daily P&L attribution	Trade reversal process
Client Data	Encryption, access controls	Weekly access logs review	Breach response plan
Regulatory	Policy templates, training	Monthly compliance scans	Violation remediation
Operational	Redundant systems	Real-time monitoring	Disaster recovery
Third-party	Vendor due diligence	Quarterly vendor reviews	Contract termination

Table 6: 90-Day NIST AI RMF Implementation Plan This project management table outlines a practical implementation timeline across 13 weeks, organized by NIST function (Govern, Map, Measure, Manage, Review). For each phase, it specifies deliverables and time requirements, totaling 60 hours over 90 days for comprehensive AI risk management implementation.

Table 6. 90-Day NIST AI RMF Implementation Plan.

Week	NIST Function	Deliverables	Time Required
1-2	Govern	AI Use Policy, Role assignment	8 hours
3-4	Map	Risk register, Impact assessment	12 hours
5-8	Measure	Metrics dashboard, Testing plan	20 hours
9-12	Manage	Controls implementation, Training	16 hours
13	Review	Full framework review, Gap analysis	4 hours
Total			60 hours

14.5. Regulatory Compliance Tables

Table 7: Investment Advisers Act Compliance for Agentic Systems This regulatory adaptation table maps traditional Investment Advisers Act requirements to agentic AI implementations. For each regulatory requirement (Fiduciary Duty, Suitability, Best Execution, Full Disclosure, Books & Records), it contrasts traditional compliance approaches with agentic AI adaptations, providing a compliance roadmap for AI-enhanced RIAs.

Table 8: State Registration Considerations for Agentic RIAs This state-specific compliance table analyzes AI-related requirements across five key states (California, New York, Texas, Florida, Illinois). For each state, it identifies AI-specific requirements, provides agentic compliance strategies, and notes relevant exemptions, enabling multi-state operations planning.

Table 9: Explainability Requirements for AI Investment Systems This technical compliance table categorizes explainability requirements across five regulatory contexts (Client Communication, Compliance Review, Examination Response, Dispute Resolution, Risk Assessment). For each context, it specifies the explainability requirement, technical implementation approach, and implementation difficulty level.

Table 10: Required AI System Documentation This documentation requirements table outlines six essential document types (AI Use Policy, Model Documentation, Testing Records, Incident Logs, Training Records, Client Disclosures) with content requirements, review frequencies, and retention periods, providing a complete documentation framework for regulatory compliance.

Table 7. Investment Advisers Act Compliance for Agentic Systems.

Regulatory Requirement	Traditional Compliance	Agentic AI Adaptation
Fiduciary Duty	Act in client's best interest	AI agents must be programmed with fiduciary constraints [12]
Suitability	Recommendations suitable for client	AI profiling must consider all client-specific factors [39]
Best Execution	Seek best execution for trades	AI algorithms must optimize execution across venues
Full Disclosure	Disclose all material facts	Disclose AI usage, limitations, and conflicts [32]
Books & Records	Maintain required records	Automated logging of all AI decisions [35]

Table 8. State Registration Considerations for Agentic RIAs

State	AI-Specific Requirements	Agentic Compliance Strategy	Exemptions
California	Data privacy requirements (CCPA)	Data localization for California clients	None for RIAs
New York	Cybersecurity requirements (23 NYCRR 500)	Enhanced AI system security controls	Small firm modifications
Texas	Disclosure of algorithmic methods	Transparent AI methodology documentation	Manual review option
Florida	Third-party vendor oversight	Rigorous AI vendor due diligence	Limited to certain AUM
Illinois	Biometric data protection	No facial recognition in client verification	BIPA compliance

Table 9. Explainability Requirements for AI Investment Systems.

Regulatory Context	Explainability Requirement	Technical Implementation	Difficulty
Client Communication	Simple explanation of AI recommendations	Natural language summaries	Low
Compliance Review	Detailed decision rationale	Decision trees, confidence scores	Medium
Examination Response	Complete audit trail	Full logging with timestamps	High
Dispute Resolution	Transparent decision factors	Feature importance analysis	Medium
Risk Assessment	Model limitations and assumptions	Model cards, datasheets	Low

Table 10. Required AI System Documentation.

Document	Content Requirements	Frequency	Retention
AI Use Policy	Permitted uses, restrictions, oversight	Annual review	5 years
Model Documentation	Architecture, data sources, limitations	Model changes	Life of model
Testing Records	Validation results, performance metrics	Quarterly	5 years
Incident Logs	AI errors, interventions, resolutions	Real-time	5 years
Training Records	Staff AI competency training	Annually	5 years
Client Disclosures	AI usage explanation, risks, benefits	Account opening	5 years

14.6. Significance and Integration

Collectively, these figures and tables provide a comprehensive visual and analytical framework for understanding, implementing, and managing agentic investment firms. They serve multiple purposes:

- **Architectural Guidance:** Figures 1-7 provide blueprints for technical implementation
- **Risk Management:** Figures 8-9 and Tables 3-6 enable effective AI governance
- **Financial Analysis:** Figure 6 and Table 2 demonstrate economic viability
- **Regulatory Compliance:** Figures 5, 10 and Tables 7-10 ensure regulatory adherence
- **Practical Implementation:** All elements combine to provide actionable guidance for small teams

The integration of these visual and tabular elements throughout the paper creates a cohesive narrative that bridges theoretical concepts with practical implementation, making the agentic investment firm model accessible and actionable for small RIAs and boutique investment teams.

15. Conclusions

The emergence of the agentic investment firm model represents a profound structural shift in the investment management landscape, fundamentally altering the competitive dynamics between small teams and large institutions. Through the strategic integration of Generative AI and Agentic AI systems, this paper has demonstrated that small RIAs and boutique firms can transcend traditional resource constraints to achieve institutional-grade capabilities with minimal overhead.

Our comprehensive analysis establishes several critical contributions to both theory and practice. First, we have presented a scalable, layered technical architecture specifically designed for resource-constrained teams, with specialized AI agents for due diligence, macro intelligence, compliance automation, and portfolio management. This architecture, complemented by a pragmatic 16-week implementation roadmap, provides a clear pathway for small firms to operationalize agentic capabilities while maintaining focus on core investment activities.

Second, we have developed a governance and risk management framework that adapts sophisticated standards like the NIST AI RMF to the realities of small-team operations. By emphasizing proportional, risk-based approaches rather than burdensome bureaucratic structures, we enable 3-person RIAs to implement robust AI governance with approximately 60 hours of effort over 90 days. This framework ensures that small firms can safely leverage autonomous systems while maintaining regulatory compliance and client trust.

Third, our detailed analysis of securities regulations—spanning the Investment Advisers Act of 1940, SEC guidance, and state Blue Sky Laws—provides a crucial bridge between technological innovation and regulatory compliance. We demonstrate how agentic systems can not only operate within existing regulatory frameworks but can be designed to enhance compliance through automated monitoring, documentation, and reporting, creating a virtuous cycle of technological advancement and regulatory assurance.

The economic implications are transformative: operational cost reductions of 50-70%, analytical bandwidth increases of 300-500%, and the ability to deliver hyper-personalized portfolio management previously accessible only to ultra-high-net-worth clients of large institutions. Critically, this model decouples analytical capability from human headcount, enabling small teams to manage “big money” through intelligent force multiplication rather than brute-force scaling.

Looking forward, the agentic investment firm model presents both challenges and opportunities. Technical challenges around model explainability, data privacy, and integration complexity must be addressed through continued innovation. Regulatory frameworks will continue to evolve, requiring agile adaptation by implementing firms. However, the foundational work presented here provides a robust starting point for small teams to navigate this evolving landscape.

The future of investment management will not be determined solely by asset size or human headcount, but by the strategic integration of human expertise with autonomous AI agency. Small, agile firms that embrace this agentic model stand to gain significant competitive advantages—not by mimicking large institutions, but by leveraging technology to enhance their inherent strengths of focus, adaptability, and client-centric service. As AI capabilities continue to advance, the agentic investment firm represents not merely an incremental improvement, but a fundamental reimagining of how investment management can and should be practiced in the 21st century.

Conflicts of Interest: The views expressed are those of the author and do not represent any affiliated institutions. This work is conducted as part of independent research. This is a review paper, and all results, proposals, and findings are derived from the cited literature. The author does not claim any novel findings. The author's work was to review and organize existing research. The LaTeX code in this paper was developed with the assistance of GitHub Copilot and edited through DeepSeek. This is a working paper and edits are expected in the next version.

References

1. Aladdin® by BlackRock - software for portfolio management. BlackRock. [Online]. Available: <https://www.blackrock.com/aladdin>
2. S. M. Bartram, J. Branke, and M. Motahari. Artificial Intelligence in Asset Management. [Online]. Available: <https://papers.ssrn.com/abstract=3510343>
3. National Institute of Standards and Technology (US), "Artificial intelligence risk management framework : Generative artificial intelligence profile," pp. error: 600–1. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
4. S. Joshi, "Comprehensive Review of Artificial General Intelligence AGI, Agentic AI and GenAI: Current Trends and Future Directions," vol. 6, no. 3, pp. 681–688. [Online]. Available: <https://www.allmultidisciplinaryjournal.com/search?q=MGE-2025-3-124&search=search>
5. Agentic AI vs. Generative AI | IBM. [Online]. Available: <https://www.ibm.com/think/topics/agentic-ai-vs-generative-ai>
6. Satyadhar Joshi, "Advancing innovation in financial stability: A comprehensive review of ai agent frameworks, challenges and applications," vol. 14, no. 2, pp. 117–126. [Online]. Available: <https://journalwjaets.com/node/310>
7. S. Joshi, "Review of autonomous systems and collaborative AI agent frameworks," vol. 14, no. 2, pp. 961–972. [Online]. Available: <https://journalijsra.com/content/review-autonomous-systems-and-collaborative-ai-agent-frameworks>
8. How Agentic AI will transform financial services. World Economic Forum. [Online]. Available: <https://www.weforum.org/stories/2024/12/agentic-ai-financial-services-autonomy-efficiency-and-inclusion/>
9. Building Effective AI Agents. [Online]. Available: <https://www.anthropic.com/engineering/building-effective-agents>
10. ByteByteGo. How OpenAI, Gemini, and Claude Use Agents to Power Deep Research. [Online]. Available: <https://blog.bytebytego.com/p/how-openai-gemini-and-claude-use>
11. BNY, Goldman double down on tech with 'digital employees' and AI agents. [Online]. Available: <https://www.investmentnews.com/wirehouses/bny-goldman-double-down-on-tech-with-digital-employees-and-ai-agents/265198>
12. AI Compliance for RIAs: Key Risks and Best Practices. [Online]. Available: <https://www.ncontracts.com/nsight-blog/investment-advisers-artificial-intelligence>
13. Here's how RIAs and tech vendors think the SEC will regulate AI. Pro Buyer. [Online]. Available: <https://citywire.com/pro-buyer/news/heres-how-rias-and-tech-vendors-think-the-sec-will-regulate-ai/a2466694>
14. Why Every RIA Needs an AI Use Policy. Moeller Law PLLC. [Online]. Available: <https://www.moellerbusinesslaw.com/blog/why-every-ria-needs-an-ai-use-policy>
15. M. Patel and J. Antonakakis. Artificial Intelligence for RIAs: Compliance Considerations and Practical Policies. Stark & Stark PC. [Online]. Available: <https://www.stark-stark.com/news/artificial-intelligence-for-rias-compliance-considerations-and-practical-policies/>
16. B. Shneiderman, "Bridging the Gap Between Ethics and Practice: Guidelines for Reliable, Safe, and Trustworthy Human-centered AI Systems," vol. 10, no. 4, pp. 26:1–26:31. [Online]. Available: <https://dl.acm.org/doi/10.1145/3419764>
17. Bridging NIST's AI Risk Management Framework with Microsoft Purview Data Security. Severian. [Online]. Available: <https://severian.ghost.io/bridging-nists-ai-risk-management-framework-with-microsoft-purview-data-security/>
18. Harnessing Generative AI in Private Equity. Bain. [Online]. Available: <https://www.bain.com/insights/harnessing-generative-ai-global-private-equity-report-2024/>
19. AI-powered assistants for investment research with multi-modal data: An application of Amazon Bedrock Agents | AWS Machine Learning Blog. [Online]. Available: <https://aws.amazon.com/blogs/>

- machine-learning/ai-powered-assistants-for-investment-research-with-multi-modal-data-an-application-of-amazon-bedrock-agents/
20. Claude for Financial Services. [Online]. Available: <https://www.anthropic.com/news/claude-for-financial-services>
 21. How Investment Teams Are Leveraging Generative AI to Enhance Research. [Online]. Available: <https://www.alpha-sense.com/blog/product/investment-teams-leverage-generative-ai/>
 22. PYMNTS. Rogo Raises \$50 Million to Develop AI-Powered Investment Banker. PYMNTS.com. [Online]. Available: <https://www.pymnts.com/news/investment-tracker/2025/rogo-raises-50-million-dollars-develop-ai-powered-investment-banker/>
 23. Generative AI in Systematic Investing: The Sizzle and the Steak | Acadian Asset Management. [Online]. Available: <https://www.acadian-asset.com/investment-insights/systematic-methods/generative-ai-in-systematic-investing-the-sizzle-and-the-steak>
 24. S. Joshi, "Review of Gen AI Models for Financial Risk Management," vol. 11, no. 1, pp. 709–723. [Online]. Available: <https://ijsrcseit.com/index.php/home/article/view/CSEIT2511114>
 25. How to assess the potential impact of AI on your portfolio | J.P. Morgan Private Bank U.S. [Online]. Available: <https://privatebank.jpmorgan.com/nam/en/services/portfolio-management/sustainable-investing/insights/how-to-assess-the-potential-impact-of-ai-on-your-portfolio>
 26. AI in Asset Management: Key Trends and Outlook for 2025. [Online]. Available: <https://www.alpha-sense.com/blog/trends/ai-in-asset-management/>
 27. A. Colson. AI agents vs. RPA: A guide for accountants. Tax & Accounting Blog Posts by Thomson Reuters. [Online]. Available: <https://tax.thomsonreuters.com/blog/ai-agents-versus-rpa-a-guide-for-accountants-tri/>
 28. AI in Portfolio Management: Use Cases & Benefits [2025 Guide]. [Online]. Available: <https://acropolium.com/blog/employing-ai-for-portfolio-management-use-cases-solutions-case-studies/>
 29. Artificial intelligence for real solutions. AboutSchwab. [Online]. Available: <https://www.aboutschwab.com/mss/story/artificial-intelligence-for-real-solutions>
 30. L. Towner. AI, ML, and Automation in Asset Management: The Roadblocks to Implementation. FF News | Fintech Finance. [Online]. Available: <https://ffnews.com/thought-leader/ai-ml-and-automation-in-asset-management-the-roadblocks-to-implementation/>
 31. Generative AI in Regulatory Affairs – Five Use Cases That Are Already Real | LinkedIn. [Online]. Available: <https://www.linkedin.com/pulse/generative-ai-regulatory-affairs-five-use-cases-already-real-wakje/>
 32. Avoid The Pitfalls Of 'AI-Washing' In RIA Marketing | Gregory. [Online]. Available: <https://gregoryagency.com/news-insights/ai-washing-ria-marketing/>
 33. SEC.gov | SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence. [Online]. Available: <https://www.sec.gov/newsroom/press-releases/2024-36>
 34. How RIAs can prepare for AI-related sweep exams. RIA. [Online]. Available: <https://citywire.com/ria/news/how-rias-can-prepare-for-ai-related-sweep-exams/a2441628>
 35. Client Alert: Generative Artificial Intelligence in Financial Services: A Practical Compliance Playbook for 2026 - Shumaker, Loop & Kendrick, LLP. [Online]. Available: <https://www.shumaker.com/insight/client-alert-generative-artificial-intelligence-in-financial-services-a-practical-compliance-playbook-for-2026/>
 36. AI in Wealth Management: Use Cases, Tools, and Guidelines. [Online]. Available: <https://www.itransition.com/ai/wealth-management>
 37. C. Künzle. AI In Portfolio Management - 3 Ways How Wealth Management Banks Benefit. ZHAW Wealth Management Blog. [Online]. Available: <https://blog.zhaw.ch/wealth-management/2024/04/10/3-ways-how-wealth-management-banks-benefit-from-ai-in-portfolio-management/>
 38. S. Joshi, "Advancing Financial Risk Modeling: Vasicek Framework Enhanced by Agentic Generative AI by Satyadhar Joshi," vol. Volume 7. [Online]. Available: https://www.irjmets.com/uploadedfiles/paper/issue_1_january_2025/66819/final/fin_irjmets1738056751.pdf
 39. The incredible potential of artificial intelligence in wealth management. institutional.fidelity.com. [Online]. Available: <https://institutionalsc.trafficmanager.net/advisors/insights/topics/running-your-business/the-incredible-potential-of-artificial-intelligence-in-wealth-management>
 40. European Parliament. Directorate General for Parliamentary Research Services., *The Impact of the General Data Protection Regulation on Artificial Intelligence*. Publications Office. [Online]. Available: <https://data.europa.eu/doi/10.2861/293>
 41. S. Joshi. Compensating for the Risks and Weaknesses of AI/ML Models in Finance. [Online]. Available: <https://www.preprints.org/manuscript/202503.2245>

42. AI Act | Shaping Europe's digital future. [Online]. Available: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
43. N. Rangone and L. Megale, "Risks Without Rights? The EU AI Act's Approach to AI in Law and Rule-Making," vol. 16, no. 3, pp. 1082–1097. [Online]. Available: <https://www.cambridge.org/core/journals/european-journal-of-risk-regulation/article/risks-without-rights-the-eu-ai-acts-approach-to-ai-in-law-and-rulemaking/3AD4822C291C6591BAFD26524CD44C12>
44. Airia AI Platform | Build, Deploy & Scale Enterprise AI. Airia. [Online]. Available: <https://airia.com/ai-platform/>
45. S. Raza, R. Sapkota, M. Karkee, and C. Emmanouilidis. Responsible Agentic Reasoning and AI Agents: A Critical Survey. [Online]. Available: <https://www.techrxiv.org/users/574774/articles/1329333-responsible-agentic-reasoning-and-ai-agents-a-critical-survey?commit=2b4aa8d87387b44274d49f053425b0fc6f718838>
46. Integrating NIST AI RMF into Modern Risk Frameworks. GovWhitePapers. [Online]. Available: <https://govwhitepapers.com/whitepapers/integrating-nist-ai-rmf-into-modern-risk-frameworks/>
47. Real-world gen AI use cases from the world's leading organizations. Google Cloud Blog. [Online]. Available: <https://cloud.google.com/transform/101-real-world-generative-ai-use-cases-from-industry-leaders>
48. Bridging the ai skills gap workforce training for financial services | International Journal of Innovative Science and Research Technology. [Online]. Available: <https://www.ijisrt.com/bridging-the-ai-skills-gap-workforce-training-for-financial-services>

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.