

AI & Investment Due Diligence

A Practical Guide for Asset Managers

Essential Reading for:

Chief Operating Officers • Investor Relations Directors
Operations Managers • Compliance Officers

BackPro AI
December 2025

backpro.ai

Contents

1	Executive Summary	2
2	The Operational Challenge	2
2.1	Growing Documentation Burden	2
2.2	Resource Constraints	2
2.3	Quality and Compliance Risks	3
3	How AI Approaches Document Automation	3
3.1	Retrieval-Augmented Generation (RAG)	3
3.2	Private Deployment Architecture	3
3.3	Accuracy Verification	3
4	Governance Framework	4
4.1	Approval Workflows	4
4.2	Data Classification	4
4.3	Audit and Compliance	4
5	Practical Use Cases	5
5.1	DDQ Response Automation	5
5.2	RFP Acceleration	5
5.3	Consultant Database Maintenance	5
5.4	Investor Reporting	5
6	Implementation Approach	6
6.1	Phase 1: Pilot (4-6 Weeks)	6
6.2	Phase 2: Rollout (6-8 Weeks)	6
6.3	Phase 3: Optimization (Ongoing)	6
7	Operational Due Diligence Considerations	7
7.1	Data Security and Sovereignty	7
7.2	Accuracy and Reliability	7
7.3	Governance and Compliance	7
7.4	Business Continuity	7
7.5	Integration and Usability	7
8	Regulatory Considerations	8
8.1	Australian Context	8
8.2	Global Considerations	8
9	Key Takeaways	8
10	Questions to Ask Vendors	8
11	Conclusion	9

1 Executive Summary

Institutional asset managers face increasing pressure to scale investor relations capacity while maintaining accuracy and compliance standards. Due diligence questionnaires (DDQs), requests for proposals (RFPs), and investor reporting consume significant operational resources, with individual DDQs often requiring 12-20 hours of manual work.

Artificial intelligence technologies, specifically those based on retrieval-augmented generation (RAG), offer a pathway to automate document-intensive processes while maintaining institutional quality standards. However, implementation requires careful consideration of data sovereignty, accuracy verification, and governance controls.

This guide examines:

- Current challenges in investor relations operations
- How RAG-based AI systems work in practice
- Governance frameworks for AI in investor communications
- Practical use cases and implementation approaches
- Operational due diligence considerations for AI deployment

2 The Operational Challenge

2.1 Growing Documentation Burden

Asset managers experience consistent growth in documentation requirements:

- **DDQ Volume:** Institutional allocators increasingly require quarterly or semi-annual DDQ updates rather than annual submissions
- **Consultant Database Maintenance:** Platforms like eVestment, Mercer, and Preqin require regular updates to maintain accuracy
- **RFP Responses:** Competitive mandates often have 2-3 week turnaround requirements
- **Investor Reporting:** Monthly and quarterly reports require consistent performance commentary and attribution analysis

2.2 Resource Constraints

Investor relations teams face structural capacity limitations:

- Senior IR professionals are difficult to recruit in Australia's competitive market
- Manual processes don't scale linearly with fund growth
- Copy-paste workflows introduce version control and consistency risks
- Time spent on documentation reduces time available for relationship management

2.3 Quality and Compliance Risks

Manual documentation processes create several risk vectors:

- **Inconsistency:** Different analysts may answer the same question differently across submissions
- **Outdated Information:** Policy documents or team information may be out of date
- **Version Control:** Uncertainty about which version of a document represents current policy
- **Audit Trail:** Difficulty demonstrating what information was provided to which investors

3 How AI Approaches Document Automation

3.1 Retrieval-Augmented Generation (RAG)

RAG systems combine document retrieval with language generation to produce accurate, sourced responses. Unlike generic AI chatbots, RAG systems:

1. **Index Fund Documents:** Investment mandates, policies, factsheets, past DDQ responses, and regulatory filings are processed and indexed
2. **Retrieve Relevant Content:** When a question is asked, the system identifies the most relevant source documents
3. **Generate Responses:** The AI drafts a response using only the retrieved source material
4. **Provide Citations:** Every statement includes references to specific source documents

3.2 Private Deployment Architecture

For institutional asset managers, data sovereignty is non-negotiable. Effective AI systems must:

- Deploy within the fund's cloud environment (AWS, Azure, Google Cloud)
- Ensure no fund data is sent to public AI services
- Maintain complete audit logs of all queries and responses
- Support air-gapped deployment for highly sensitive environments

3.3 Accuracy Verification

RAG systems differ from general-purpose AI in their approach to accuracy:

- **Source Attribution:** Every generated statement references specific documents
- **Confidence Scoring:** System indicates when source material is ambiguous
- **Human Review Workflow:** Responses are flagged for review before sending
- **Version Control:** System tracks which document versions were used for each response

4 Governance Framework

4.1 Approval Workflows

Appropriate governance requires multi-layer review:

1. **AI Draft:** System generates initial response with source citations
2. **Analyst Review:** IR analyst verifies accuracy and completeness
3. **Senior Review:** Senior IR or portfolio manager reviews material answers
4. **Compliance Check:** Compliance officer reviews regulatory or risk-related content
5. **Final Approval:** Authorized signatory approves before sending

4.2 Data Classification

Not all content is equally suitable for AI processing:

Suitable for AI Processing

- Investment philosophy and process descriptions
- Organisational structure and team biographies
- Historical performance data (already public)
- Policy documents (risk, ESG, compliance)
- Standard operational procedures

Requires Additional Controls

- Forward-looking statements or performance projections
- Non-public material information
- Investor-specific terms or side letters
- Regulatory filings before public release
- Competitive positioning or strategy

4.3 Audit and Compliance

Institutional-grade AI systems must support comprehensive auditability:

- **Query Logs:** Complete record of all questions asked and responses generated
- **Source Tracking:** Which documents were used for each response
- **Version History:** Changes made during human review process
- **User Attribution:** Who initiated requests and who approved responses
- **Retention Policies:** Compliance with record-keeping requirements

5 Practical Use Cases

5.1 DDQ Response Automation

DDQs are highly suitable for AI assistance due to their structured, repetitive nature:

Process Flow:

1. Upload new DDQ to system
2. AI identifies question types and retrieves relevant past responses
3. System generates draft answers with source citations
4. Analyst reviews for accuracy and updates where necessary
5. Senior reviewer approves material sections
6. Compliance reviews risk and regulatory content
7. Final DDQ exported for submission

Time Reduction: First-draft generation can reduce initial response time by 60-70%, allowing analysts to focus on verification and customization rather than searching for past answers.

5.2 RFP Acceleration

RFP responses benefit from AI's ability to adapt past content to new formats:

Capabilities:

- Extract relevant content from past winning proposals
- Adapt historical responses to match new RFP structure
- Identify gaps where new content is needed
- Maintain consistent brand voice and messaging
- Flag potentially conflicting statements across sections

5.3 Consultant Database Maintenance

Platforms like eVestment, Mercer, and Preqin require regular data updates:

AI Applications:

- Auto-populate database fields from current factsheets
- Flag outdated information requiring updates
- Ensure consistency across multiple platforms
- Generate change logs for compliance review

5.4 Investor Reporting

Quarterly and monthly reports can leverage AI for consistent commentary:

Use Cases:

- Generate market commentary based on approved templates
- Create performance attribution explanations
- Draft portfolio updates using current holdings data
- Maintain consistent tone across investor communications

6 Implementation Approach

6.1 Phase 1: Pilot (4-6 Weeks)

Begin with controlled testing on historical documents:

- Select 5-10 completed DDQs from past 12 months
- Ingest relevant source documents (policies, factsheets, past responses)
- Process DDQs through AI system
- Compare AI-generated responses to actual submitted responses
- Measure time saved and accuracy metrics
- Identify areas requiring additional training or source material

Success Criteria:

- 80%+ of responses require only minor edits
- Zero material inaccuracies in generated content
- 50%+ reduction in initial draft time
- Positive feedback from reviewing analysts

6.2 Phase 2: Rollout (6-8 Weeks)

Expand to full team and additional use cases:

- Train entire IR team on system usage
- Establish formal review and approval workflows
- Integrate with document management systems
- Expand source document library
- Begin RFP and reporting use cases
- Implement compliance monitoring dashboards

6.3 Phase 3: Optimization (Ongoing)

Continuous improvement of system performance:

- Quarterly governance reviews with compliance
- Expand response library with approved content
- Refine AI prompts based on analyst feedback
- Integration with additional platforms
- Performance monitoring and quality metrics

7 Operational Due Diligence Considerations

When evaluating AI vendors, operational due diligence should cover:

7.1 Data Security and Sovereignty

- Where is data stored and processed?
- Is deployment in fund's own cloud environment supported?
- What encryption standards are used?
- Can the system operate in an air-gapped environment?
- How is access controlled and monitored?

7.2 Accuracy and Reliability

- How does the system handle ambiguous or conflicting information?
- What mechanisms prevent hallucination or fabrication?
- How are sources cited and verified?
- What testing has been performed on accuracy rates?
- How does the system handle questions outside its knowledge base?

7.3 Governance and Compliance

- What audit trails are maintained?
- How long are logs retained?
- Can the system demonstrate what information was provided to specific investors?
- What approval workflows are supported?
- How are regulatory changes managed?

7.4 Business Continuity

- What is system uptime and availability?
- What disaster recovery procedures are in place?
- Can the system be backed up and restored?
- What happens if the vendor relationship ends?
- Is there vendor lock-in or can data be exported?

7.5 Integration and Usability

- How does the system integrate with existing document management?
- What file formats are supported?
- How much technical expertise is required for operation?
- What training and support is provided?
- Can the system adapt to fund-specific workflows?

8 Regulatory Considerations

8.1 Australian Context

For Australian asset managers, several regulatory frameworks apply:

- **ASIC Regulatory Guidance:** No specific AI regulations yet, but general obligations around accuracy and record-keeping apply
- **Privacy Act:** Personal information handling requirements for staff and investor data
- **AML/CTF:** Record-keeping and audit trail requirements
- **Corporations Act:** General obligations around accurate disclosure

8.2 Global Considerations

For funds with international investors or regulatory obligations:

- **EU AI Act:** Risk-based framework for AI systems (2024 implementation)
- **SEC AI Guidance:** US regulatory focus on conflicts of interest and accuracy
- **FCA UK:** Technology governance requirements
- **GDPR:** Data processing and privacy requirements for EU investors

9 Key Takeaways

Core Principles

1. **Data Sovereignty:** Deploy AI in your own environment, never send fund data to public AI services
2. **Source Attribution:** Every AI-generated statement must reference specific source documents
3. **Human Oversight:** AI assists, but humans review, verify, and approve
4. **Audit Trail:** Comprehensive logging of queries, responses, and approvals
5. **Governance Framework:** Multi-layer review process before any content is sent externally
6. **Continuous Monitoring:** Regular compliance reviews and quality assessments

10 Questions to Ask Vendors

Before selecting an AI vendor, ensure you can answer:

1. Can you demonstrate actual accuracy rates on real DDQs, not synthetic examples?
2. What specific controls prevent the AI from fabricating information?
3. How do you handle questions for which there is no source material?
4. Can we deploy in our own cloud environment without data leaving our systems?

5. What audit trails and compliance reports are available?
6. How do you handle regulatory changes and policy updates?
7. What disaster recovery and business continuity provisions exist?
8. Can you provide references from similar institutional clients?
9. What is your roadmap for handling forward-looking statements and sensitive content?
10. How do we retain our data if we terminate the relationship?

11 Conclusion

AI technology offers genuine operational leverage for institutional asset managers facing growing documentation burdens. However, successful implementation requires more than just technology deployment—it requires thoughtful governance, rigorous accuracy verification, and appropriate human oversight.

Asset managers should approach AI implementation as a strategic initiative requiring:

- Executive sponsorship and governance framework
- Phased rollout with measurable success criteria
- Comprehensive operational due diligence on vendors
- Ongoing compliance monitoring and quality assurance
- Clear policies on what AI can and cannot be used for

When implemented with appropriate controls, AI can free significant IR capacity while improving consistency, reducing errors, and maintaining comprehensive audit trails. The key is treating AI as an assistant that accelerates human expertise, not a replacement for judgment and oversight.