Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Intelligent Rule Extraction and Decision Management: An AI-Driven Architecture for Enterprise Decision Optimization

[1]Ishant Goyal, [2]Gireesh Patil, [3]Amjad Shaikh

[1]ServiceNow Inc., USA

[2]ServiceNow Inc., USA

[3]ServiceNow Inc., USA

ARTICLE INFO

Received: 16 Aug 2025

Revised: 18 Sept 2025

Accepted: 27 Sept 2025

ABSTRACT

The article describes an Intelligent Rule Extraction and Decision Management (IREDM) architecture that tackles the key problem of Decision Debt in businesses where rich business logic is languishing in unchangeable policy documents. The AI-based framework automatically derives actionable rules out of unstructured documentation and translates them into computational decision models in a fully traceable manner. A three-plane architecture separates functionality into knowledge, decision, and governance planes, resulting in a modular system with an appropriate balance between automation and suitable controls. The solution itself provides the transformative gains of consistent application of rules, verifiable compliance, expedited execution of policies, scalability to an enterprise-wide level, and continuous improvement in the form of feedback loops. The document provides comprehensive examples, such as the process of automated vendor onboarding, to depict how organizations can instantiate their institutional knowledge, transforming their static policies into dynamic, governed decision capabilities that change in response to changing business and regulatory conditions.

Keywords: Decision Intelligence, Rule Extraction, Knowledge Graphs, Governance Framework, AI Agents

1. Introduction: The Enterprise Decision Challenge

1.1 The Growing Problem of Decision Debt

The current business environment confronts any modern enterprise with unprecedented challenges in dealing with the huge ecosystem of policies and procedures that dictate critical business operations. The latest studies have found an alarming pattern of organizational knowledge being stuck in stagnant reports, which is known as the Decision Debt - that is, in the pattern whereby the gap between the written policies and the actual implementation is getting larger and larger, as time passes. This disconnection can take different forms within any industry, including financial services or healthcare, manufacturing, or public administration. Companies large and small find it difficult to derive executable decision logic out of thick policy documentation, and as a result, interpret and apply the same underlying rules differently. The situation is exponentially more complicated once organizations expand internationally, and the diversity of regulations in different regions of the world adds even more complexity to the standardization of decision-making. This has been compounded by the speed at which regulatory changes are occurring in post-crisis settings, necessitating more flexible ways to manage and implement policy than traditional manualized procedures can implement [1].

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

1.2 The Hidden Costs of Manual Policy Interpretation

The net financial and operational cost of depending on manual interpretation of policy goes well past inefficiency. Organizational knowledge workers spend significant chunks of productive time reading policies and procedures, and the error rates rise exponentially as decisions become more complex. The interpretation overhead is a considerable hidden cost that is not quantified in operational budgets to a considerable degree. The policy inconsistency generates three different organizational issues that affect the bottom-line performance and regulatory status. First, the difference in interpretation results in inconsistent customer experiences and internal operations, and significant changes in service quality may be observed between business units that follow the same policy frameworks. Second, the subjective chain of evidence linking policies with operational decisions introduces serious vulnerabilities at the level of regulatory audits and reviews of oversight. Lastly, the long time to decision caused by the manual interpretation process has a direct effect on the agility and responsiveness of organizations to market changes, which generates competitive disadvantages in agile industries [2].

1.3 The Intelligent Rule Extraction Paradigm

One of the most effective approaches to these challenges has been developed by applying sophisticated artificial intelligence methods to understanding policies and extracting rules. This smart rule extraction paradigm is a paradigm shift in the operationalization of documented policies in an organization and forms a direct connection between written procedures and executable decision logic. The architectural method embraces advanced natural language processing to infer semantic meaning out of the unstructured text, finding the elements of the decision that are decision-critical, such as conditions and actions, and formalizing the elements of the decision into structured rule representations. With clear traceability between source documents and operational rules, organizations have never had a clearer view of the origin of each business decision. The mechanisms of continuous learning that are built into these systems mean that as new policies are brought in by amendments and updates, the operational rule base automatically changes to ensure that the documentation and the execution are in alignment. The methodology provides considerable benefits as measured in various aspects such as consistency of decision-making, ability to check compliance, and efficiency in operations without compromising administrative authority and accountability [1].

2. Concept Overview

Fundamentally, the Intelligent Rule Extraction and Decision Management (IREDM) concept is an advanced cognitive system that improves the way organizations manage complex policy settings. This way represents the development of artificial intelligence technology that once was simple rule-based systems to autonomous agents who are able to see what is happening around them, make decisions, and take actions to accomplish certain aims. AI agents are digital workers that enhance human skills, but are not going to substitute them, especially in areas that involve specialized skills such as policy analysis and regulatory compliance. These agents have different levels of autonomy according to the sensitivity and complexity of the tasks they are given. Agents can be set to work with suitable guardrails in regulatory settings to be responsive to organizational governance needs, yet still provide substantial efficiency gains. This idea extends earlier technologies such as natural language processing, knowledge representation, and machine learning to develop systems that can learn the intricacies of policy language and convert it into operational decision systems. Handling this agent-based approach introduces a new dimension between written documentation and active business processes and allows firms to have a better fit between written policies and carried-out activities and respond faster to regulatory changes [5].

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

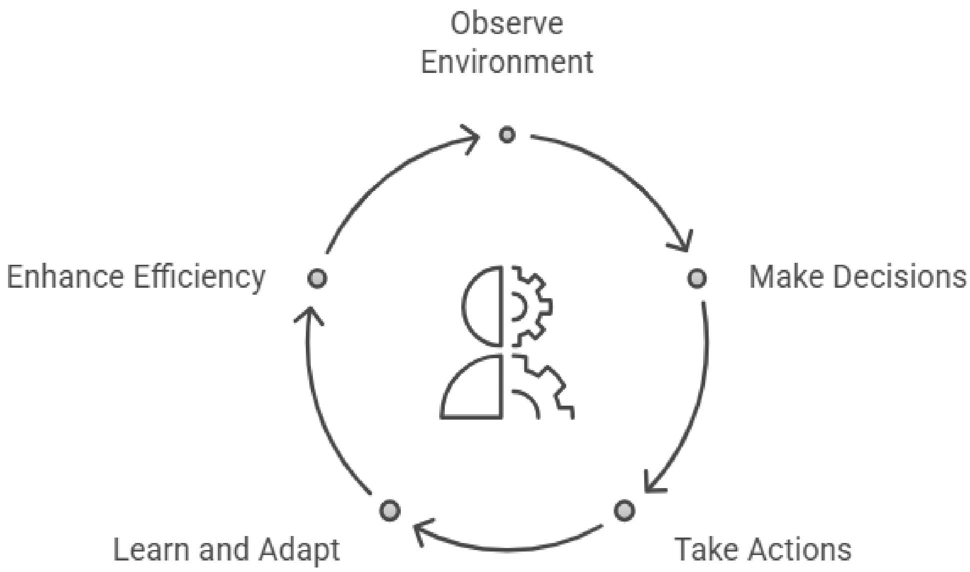

Fig 1: IREDM Cycle [4, 5]

2.1 What the AI Agent Does

The following are the main benefits of AI Agents of IREDM:

- **Improved Document Processing:** Processes multiple financial and regulatory documents of any format, deriving valuable information and eliminating unnecessary information automatically [5].
- **Intelligent Rule Extraction:** Recognizes sophisticated decision criteria found in policy language that contains conditional structure, numerical rules, and time limits that specify how particular situations are to be dealt with [6].
- **Formatted Knowledge Representation:** converts unstructured policy text into machine-readable form with clear provenance to source documents to support both automated decisions and audit needs [6].
- **Automated Decision Execution:** Compares business situations to extracted business rules and decides what should be done, and gives reasons that meet operational and regulatory transparency criteria [6].
- **Exception Handling**: Recognizes unusual or risky cases that involve human intervention within configurable thresholds that provide a balance between automation and reasonable supervision [6].
- **Continuous Improvement**: Adds feedback to optimize models based on feedback and evolves with policy changes, with automatic documentation update detection [5].
- **Reduced Processing Time**: Decelerates document-intensive business operations such as loan origination by automatically executing rules application and verification procedures [6].
- **Better Compliance:** Enforces consistency between the operational decision and the existing regulatory requirements by applying and documenting rules systematically [5].

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

3. Architecture Overview: The IREDM Framework

3.1 Architectural Principles and Design Philosophy

The IREDM architecture uses the tripartite framework that breaks functionality into different planes. This solution is in line with the principles of integrated risk management in that it establishes distinct boundaries between the knowledge processing, decision execution, and governance functions. The segregation allows the specific areas to be improved without interfering with the whole system, whilst the entire system is highly accountable through the entire lifecycle of decision making [3].

3.2 Knowledge Plane Components

3.2.1 Policy Ingestion Layer

This base layer links to various document bases in the organisation, converting fixed content into available knowledge bases. It includes the ability to acquire documents, processing pipelines to clean them, and semantic indexing to retrieve documents based on meaning. These elements form a single content base that can be used in further extraction processes and that maintains links to the authoritative materials [3].

3.2.2 Policy Understanding & Rule Extraction Layer

This layer is the cognitive core of IREDM and is used to convert unformatted policy text to structured business rules using domain-relevant language models. It understands terms in the industry, finds important objects in policies, and represents complex conditional and temporal logic. This domain expertise represents the segment architecture strategy outlined in enterprise architecture models, which generates active knowledge resources out of passive documentation [4].

3.2.3 Knowledge Graph & Rule Repository

This layer provides a graph-based structure, which encodes the individual rule, as well as the interrelationships among the rules, in a computational indexable format. The approach encapsulates the inherent associations among organizational policies and attaches holistic metadata such as provenance, applicability parameter, and governance features. The resulting repository offers a single-centralized view of organizational decision logic that can be easily traced to original documents [3].

3.3 Decision Plane Components

3.3.1 Decision Engine

The Decision Engine compares real-time business conditions with the rule base extracted to identify the right actions. It implements various reasoning patterns of simple deterministic rules to elaborate chains of inferences and gives a score of confidence to each output. The extent to which this component represents the service component reference model in enterprise architecture is that it represents business logic in the form of reusable services that can be called upon to perform a variety of business processes [4].

3.3.2 Workflow Integration

This layer links the IREDM framework to operational systems by standardized integration patterns and APIs. It converts the result of a decision into tangible business operations and handles exceptions by intelligently routing them to special handling queues. These are capabilities that are in line with the model of the technical reference of the enterprise architecture frameworks, where standard approaches to interconnection between systems are promoted to ensure uniformity across business processes [4].

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

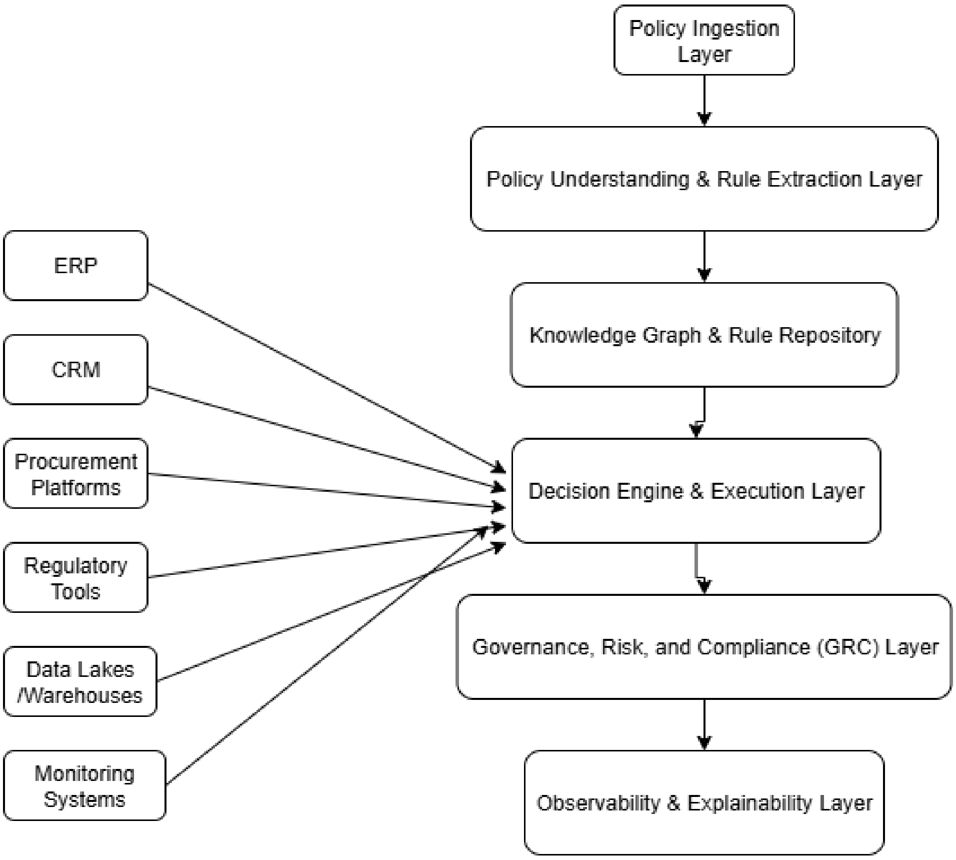

Fig 2: Scalable and Governable IREDM System Architecture [3, 4, 6]

3.4 Governance Plane Components

3.4.1 GRC Layer

The Governance, Risk, and Compliance layer provides controls that keep the system within the right limits. It includes compliance guardrails, human-in-the-loop orchestration to enforce the right oversight, and formal rule validation workflows. These elements contribute directly to the integrated risk management through a systematic approach to mitigation, risk assessment, and risk identification during the decision lifecycle [3].

3.4.2 Observability & Explainability Layer

This layer offers critical insights into the decision operations with a full accounting of the operations, explainable AI services, and role-based performance dashboards. It puts the entire context of every decision in place and produces human-readable explanations of results. These features allow the performance reference model in the enterprise architecture to define concise measures and continuous enhancement of operations through feedback [4].

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

4. Agentic Workflow Example: Automated Vendor Onboarding

Ingestion & Understanding

Onboarding of a vendor starts with the AI agent reviewing key procurement policy and compliance documents. The agent uses document understanding technology to extract unstructured data in various formats to identify parts, hierarchical relationships, and cross-references among parts of a policy. This contextual processing constructs a holistic map of the requirements related to vendor management upon which rule extraction and decision automation are based [7].

Extraction

The agent determines particular vendor management needs that are incorporated in the policy documentation and targets specific actionable business rules that have a direct effect on qualification and approval processes. These most often involve financial limits, geographical limitations, certification, and procedural rules. The extraction uses named entity recognition to find certain values, timeframes, roles, and compliance criteria, and generates a hierarchical depiction of the vendor governance framework within the organization [7].

Structuring

The rules obtained are presented in a knowledge graph of relationships between and among individual requirements. Attributes in each rule node comprise the statement in its exact form, classification, category of vendor that the rule applies to, and source information. The graph topology allows traversal functions to determine all the applicable rules in a particular situation, and there is easy tracing of the rule to the authoritative source, translating the formal policy description into an implementation decision model [8].

Decision Execution

The agent compares the submission to the structured rule base when a new vendor application is introduced into the system. This involves geographic risk evaluation, financial threshold evaluation, and service category evaluation in order to determine specialized needs. The agent uses external data enrichment to increase the quality of the decision, gathering automatically company information and verification by trusted sources. In the process, full traceability is ensured between every evaluation stage and policy requirements beneath [8].

Output & Justification

The multi-part responses generated by the agent contain both action recommendations and detailed justifications associated with particular sections of policy. In the case of financial reviews, the agent produces routing instructions to the correct approval queues. In documentation requirements, it makes us start the communication processes that demand particular certifications. All components have structured annotations of references to certain sections of the policy, which makes the policy an active aspect of the working process [7].

Learning & Adaptation

The ability to detect document changes automatically detects policy updates and triggers specialized re-extraction processes. The governance controls make sure that before a modification is implemented, it must be validated. Feedback loops receive feedback about specialists who are examining agent suggestions, polishing extraction competencies as time progresses. Usage analytics can be used to see trends in submissions and approvals that point to potential policy improvements or process refinements [8].

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

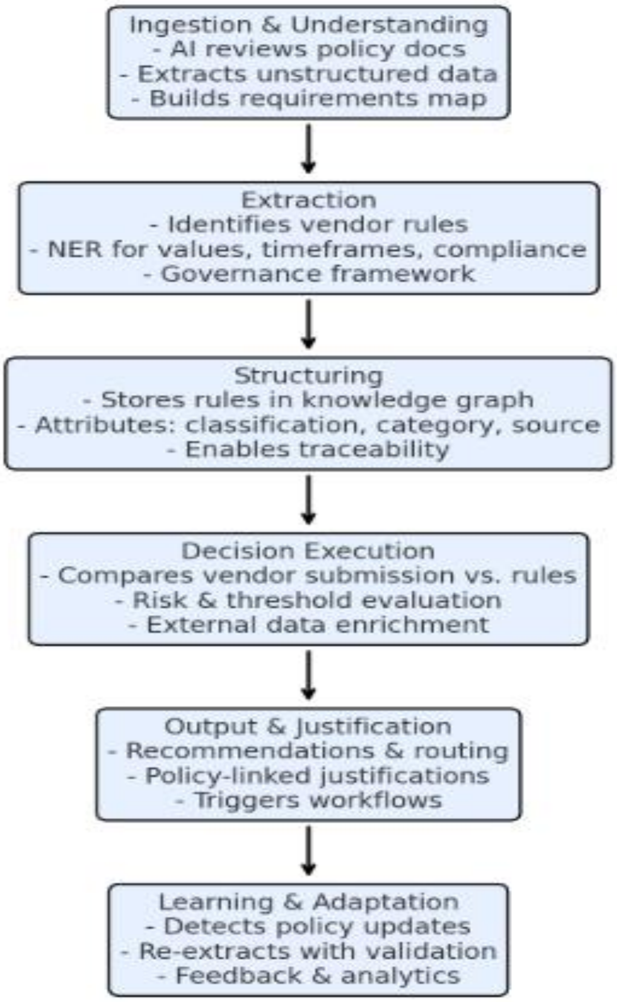

Fig 3: Automated Vendor Onboarding agentic process [6, 7, 8]

5. Governance and Control Framework

Governance of AI agents needs to be an all-inclusive process that cuts across the entire decision lifecycle. An effective governance structure balances technical and organizational components of agent deployment and provides opportunities to tackle a range of AI-related risks proactively, balancing innovation with corresponding protective measures. Such a solution keeps intelligent systems within reasonable limits and provides quantifiable business value [9].

5.1 Pre-Deployment (Design-Time) Controls

Pre-deployment governance provides a set of essential underpinnings before an AI agent processes its initial document. Policy curation systems assign official bodies, so that the agent acquires knowledge based on existing, approved records, and not obsolete ones. Bias and risk assessment procedures measure how the language models understand policy needs, and it is specifically important to consider how the automated decisions made might cause unfairness to a particular class of people or situations in which automated decisions may pose a risk of inequality. Structured review processes are part of formal rule validation in which extracted rules are reviewed by subject matter experts in business, legal, and compliance functions, such as completeness checks, accuracy checks, and consistency checks [9].

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

5.2 Runtime (In-Flight) Controls

Dynamic controls are used to offer continuous control over agents in operation. Confidence scoring attributes reliability measures to every decision, allowing intelligent routing (high-confidence decisions run through automation, other decisions run through human inspection). Tiered risk ranking is the categorisation of decisions based on potential impact and may include factors such as financial materiality, regulatory sensitivity, and reputational risk. Decisions that are at a higher level are automatically given superior scrutiny despite confidence scores. Unchangeable audit trails provide end-to-end documentation of decision processes in stored, immutable storage, and document both inputs, rules to be applied, and the logic steps and results to serve governance requirements such as compliance and performance assessment [9].

5.3 Post-Decision (Post-Facto) Controls

Post-facto controls provide essential feedback mechanisms for responsible AI governance. Continuous outcome analysis measures decision performance across multiple dimensions, including accuracy, compliance adherence, business impact, and efficiency gains. Proactive change management addresses policy evolution through automated monitoring of source repositories, systematic comparison of document versions, and controlled validation before updated rules become operational. User feedback loops enable stakeholders to question automated determinations, serving dual purposes: addressing specific decision concerns while generating valuable training data that improves system performance over time [10].

6. Monitoring, Explainability, and Observability (MEO)

Enterprise AI implementation is based on transparency and understanding. The Monitoring, Explainability, and Observability (MEO) layer develops various complementary systems that reveal the internal processes of an AI-based decision system to establish stakeholder trust and meet regulatory demands of algorithmic accountability [11].

Decision provenance and lineage functions are the building blocks of trust as they establish a two-way traceability between all decision components and their authoritative origin in policy records. Stakeholders question a determination and get immediate access to the policy clause, sentence, and page that informed the determination. This traceability encompasses timestamps, versioning, and approval data that form a full chain of custody on the decision logic and which provides an auditable relationship between operational choices and governance records [11].

Multi-layered explainability turns black box AI decisions into transparent and comprehensible procedures. Complementary methods such as feature importance analysis, visualization of attention, and counterfactual explanations are effective explainable AI methods that help users understand how the system understands policy language. Natural language explanations are valuable to business users because they encode technical justifications in understandable accounts that relate decisions to familiar business concepts. This multidimensional solution will meet technical validation requirements to regulatory justification requirements, to user acceptance requirements [11].

Continuous visibility is brought about by performance and health monitoring, which includes role-specific dashboards that monitor metrics in various dimensions. These dashboards record operational throughput, quality metrics, as well as compliance metrics with alerting systems that actively signal stakeholders when the metrics fall out of range, allowing them to respond quickly to arising problems [11].

The automated exception and conflict detection detects situations that demand special attention, such as policy conflicts where alternative rules yield conflicting guidance, ambiguous situations that have no specific guidance, and coverage gaps where the policies fail to sufficiently capture business

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/
Research Article

situations. These processes convert passive records into functioning governance instruments that can constantly assess the decision structure within the organization [11].

7. Strategic Benefits

Deployments of Intelligent Rule Extraction and Decision Management architecture provide revolutionary benefits that redefine the way organizations approach policy-based decision making, generating strategic effects larger than tactical efficiency benefits [12].

Radical consistency and fairness are achieved when policy documents are used to extract decision logic and applied systematically to all transactions. This cuts out the variation of human interpretation where the same situation is frequently treated differently by different staff members, or by the same staff member in different situations or times. Organizations that have adopted IREDM achieve large changes in fairness measures as they do not worsen the overall performance or remain unchanged. This consistency applies to customer-facing situations as well as those internal to the organisation, such as procurement, compliance reporting, and risk management [12].

Regulatory confidence that can be verified changes compliance posture in that it develops direct, verifiable interconnections between all operational choices and the mandate behind them. This traceability provides compliance by design, where compliance is incorporated into operations as opposed to being evaluated after the fact. When audits are conducted, the ability to show accurate lineage of decisions to governing policies can significantly lower findings and remediation needs by the organization. This proactive approach enables audit processes to be more of an exercise that consumes resources than a simple demonstration of systematic controls [12].

It offers competitive advantages of accelerated speed and efficiency by providing very short cycle times in creating policies and moving them into operation. Automated extraction and deployment are going to shrink the time frame, whereas the more conservative methods require months of interpretation, procedure writing, and staff training. The top implementations are capable of deploying standard policy changes in under a day and implementing complex regulatory changes in under one day. This is carried over to decision execution, where processing time is cut from days to minutes, and quality is maintained [12].

The huge scalability allows uniform decision-making across geographical locations, organizational sectors, and volumes of transactions without scaling the resources. They can use one proven set of decision logic across the enterprise and make the required local adjustments. This architecture can scale up to meet volume spikes and scale out to market new business opportunities without a corresponding linear increase in resources, achieving operational elasticity without impacting governance controls [12].

With continuous improvement mechanisms, decision intelligence is developed and reinforced as time passes by, as built-in learning cycles that encapsulate operational feedback, exception patterns, and performance metrics. It forms a virtuous circle in which every transaction adds some intelligence to the system. In addition to the operational benefits, the analytical capabilities offer strategic feedback concerning the effectiveness of the policy that can be refined, evidence-based, to bring the governance frameworks closer to organizational goals [12].

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Benefit	Key Point
Consistency & Fairness	Uniform policy application, eliminating human variation.
Regulatory Confidence	Verifiable traceability, easier audits, compliance by design.
Speed & Efficiency	Rapid policy deployment, faster decision execution.
Scalability	Handles high volumes and geographies without extra resources.
Continuous Improvement	Feedback loops enhance decision intelligence over time.

Table 1: Strategic Benefits of Intelligent Rule Extraction and Decision Management (IREDM) [12]

Conclusion

The process of replacing manual interpretation of policies with AI-based rule extraction is a significant paradigm shift in how businesses operationalize their governance systems, building a dynamic decision-making fabric that continuously changes with the needs of an organization and regulatory demands.

IREDM architecture decouples concerns on knowledge, decision, and governance planes, forming a modular model that can evolve in small steps. The knowledge plane turns the chaotic documentation into rules that are fully traceable. The decision plane can apply these rules fairly throughout and explain everything transparently. The governance plane is a balance between automation and risk management, as well as the governance of the whole decision lifecycle.

Organizations that have adopted IREDM systems report a dramatic increase in operational consistency because the same situations have been treated in the same way, irrespective of channel or staff. Verification of compliance becomes dynamic, not passive, and each decision can be traced back to policy. The pace of the process increases dramatically, and policy changes become visible in operations within days, not months. The architecture is scaled efficiently throughout the enterprise, and continuous learning mechanisms maintain an ever-increasing quality of decisions.

The IREDM will evolve in several emerging directions, such as multimodal understanding that goes beyond text and uses visual elements in documenting policies, improved human-AI interactions that can generate hybrid intelligence models, and predictive analytics to simulate possible changes in policy prior to the implementation. As AI governance regulatory frameworks evolve, IREDM architectures will include newer controls to ensure that they remain aligned with new requirements.

The path to the decision-intelligent enterprise opens up the real worth of institutional knowledge and transforms the inert policies in the enterprise into dynamic decision capabilities that provide consistent, compliant, and assuring outcomes throughout the enterprise.

Journal of Information Systems Engineering and Management
2025, 10(60s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

References

[1] Danyu Zhao et al., "Multiagent Decision-Making of Enterprise Debt Restructuring Based on Fuzzy-Set Qualitative Comparative Analysis," Wiley, 2023. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1155/2023/6722707

[2] IMF, "Revised Guidelines for Public Debt Management," 2014. [Online]. Available: https://www.imf.org/external/np/pp/eng/2014/040114.pdf

[3] Nick Barney and Wesley Chai, "What is integrated risk management (IRM)?" TechTarget, 2025. [Online]. Available: https://www.techtarget.com/searchsecurity/definition/integrated-risk-management-IRM

[4] "Federal Enterprise Architecture Framework," ObamaWhiteHouse, 2013. [Online]. Available: https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/egov_docs/fea_v2.pdf

[5] Anna Gutowska, "What are AI agents?" IBM. [Online]. Available: https://www.ibm.com/think/topics/ai-agents

[6] Kushagra Bhatnagar, "AI Agents: The Future of Document Processing in Lending," Arya.ai, 2025. [Online]. Available: https://arya.ai/blog/ai-agents-document-processing-lending

[7] Guru Angisetty, "Vendor Onboarding Agent: Automate Supplier Integration with No-Code AI," ODIN, 2025. [Online]. Available: https://blog.getodin.ai/vendor-onboarding-agent/

[8] Oracle, "AI Agents in Oracle Procurement: Automating Policy Checks, Supplier Onboarding, and More," 2025. [Online]. Available: https://www.2-data.com/knowledge-hub/ai-agents-in-oracle-procurement-automating-policy-checks-supplier-onboarding-and-more

[9] Cole Stryker, "AI agent governance: Big challenges, big opportunities," IBM. [Online]. Available: https://www.ibm.com/think/insights/ai-agent-governance

[10] Gal Golan, "AI Governance Framework: Key Principles & Best Practices," MineOS, 2025. [Online]. Available: https://www.mineos.ai/articles/ai-governance-framework

[11] IBM, "What is explainable AI?" 2023. [Online]. Available: https://www.ibm.com/think/topics/explainable-ai

[12] Mukul Bhati, "Revolutionize Your Business with Enterprise Decision Management," Nected, 2024. [Online]. Available: https://www.nected.ai/blog/enterprise-decision-management

[13] Harishchandra Patel, “Impedance Control in HDI and Substrate-Like PCBs for AI Hardware Applications” (2024). Journal of Electrical Systems, 20(11s), 5109-5115.