Governed RAG-Agent Architecture for Autonomous Enterprise Decisions

Emma Richardson
Asst Professor

Abstract

Enterprises seek to harness emerging technologies such as generative AI to enable decision automation in business operations while maintaining governance, risk, and oversight. Existing architectures for building Autonomous Agents in Information Systems work well for narrow-band decision problems but fall short when adapting generative AI methods to solve problems requiring high governance and compliance risk mitigation. Existing ServiceNow implementations operate at ChatOps levels of automation but lack the otherwise available major platform functionality. Governance-Embedded Action-Policy Decision Intelligent architecture provides an end-to-end pipeline, embedded policy enforcement, built-in audit/traceability, and governance oversight of decoupled/parameterized ServiceNow Decision Intelligent Workflows.

The concept is being explored and tested in a large enterprise with the necessary conditions for success. The initial focus is business-process auto-detection and classification, with connector integration into enterprise systems such as global knowledge repositories, tracking systems, and data quality monitoring. These endpoints provide situational updates for Business-Process Decision Intelligence reflex workflows.

Keywords: Generative Artificial Intelligence, Decision Automation, Enterprise Governance, Autonomous Agents, Compliance Risk Mitigation, Governance-Embedded Architecture, Action-Policy Decision Intelligence, ServiceNow Workflows, Embedded Policy Enforcement, Audit And Traceability, Intelligent Business Processes, Decision Intelligent Pipelines, Business-Process Auto-Detection, Process Classification Systems, Enterprise System Integration, Knowledge Repository Connectivity, Data Quality Monitoring, Situational Awareness Analytics, Reflex Decision Workflows, Scalable Enterprise AI.

1. Introduction

Enterprise applications generate vast, valuable sources of operational data—yet, manipulated manually, such information is of limited use and cannot be counted on to enhance timely decision-making nor optimize resource allocation. Recent work proposes a pipeline architecture capable of continuous autonomous intelligence creation through the intelligent orchestration of all ServiceNow workflows. Encompassing sensing, decisioning, and acting, the pipeline exploits knowledge graphs and large-language models to create new knowledge-based workflows as required for decision points. Enterprise governance is paramount: automated decision-making must balance the reach of artificial intelligence with the risk of undesirable consequences. Governance embeds itself in the architecture through policy-driven RAG–agent safeguards that wire policy and oversight directly into the pipeline. Underserved decision intelligence use cases benefit most; an enterprise organization mandated to leverage ServiceNow as a digital operating system provides the experimental context. Policies facilitate auditing, transparency, and traceability. Enterprises generate enormous amounts of data yet are notoriously slow to act on new knowledge, even though the consequences of delay can be significant, and even disastrous. Autonomously intelligent enterprise processes are rare. Still less common are decision points that are fully autonomous—that is, for which the outcome of the decision is executed in full without human input. Factors limiting the breadth and depth of automation include limited scope, narrow formality and interpretability, high evaluation costs, regulatory and vice-limiting policy requirements,

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

mitigation of high-risk outcomes, and unbounded general policies.

1.1. Key Objectives and Scope of the Study

Research and development (R&D) decisions based on data analysis often delay project rollouts, postponed systems integrations, and customer satisfaction. To embed governance, risk management, and value creation—a skeleton for an enterprise's shape and health—into enterprise intelligence, the orchestration backend ServiceNow is proposed. Using the ServiceNow Decision Intelligence Operating Model, an end-to-end orchestration, responsiveness, and agility flow is created for enterprise Decision Intelligence. Autonomous workflows with a clearly defined sense–decide–act control loop reduce response time but introduce additional risks. To address these, a Risk–Agility–Governance–Agent architecture embeds governance at the elements forming the control loop, ensuring that operational safety envelopes are defined for sensed parameters, mandates for controls removing workflows from the safety envelope, and fallback action procedures for the autonomous script.

The priorities of a post-COVID organizational structure have shifted to enhance responsiveness, speed, and agility. Enterprises want to be data-driven and sensitive to external market fluctuations while accurately sensing and responding to situations and requirements. Echoing these sentiments, the ServiceNow Decision Intelligence Operating Model proposes a single end-to-end digitally operated, orchestrated, and integrated enterprise with little or no role for humans in underlying processes. Project decisions involving data analysis, such as R&D, project implementation, and support, currently take the longest to execute. The ServiceNow platform comes equipped with the necessary set of native and added-on modules and integrates with third-party systems and applications.

Fig 1: Orchestrating Autonomous Enterprise Intelligence: A Risk–Agility–Governance–Agent Architecture for Decision-Driven R&D Workflows

2. Background and Motivation

Decision Intelligence is growing in importance across enterprises, and large-scale integration and optimization of RAG Agents are increasingly being considered for real-world use. Previous studies have identified a range of objectives for achieving Decision Intelligence, coupled with related challenges and opportunities. Building an RAG-Agent architecture on the ServiceNow platform would seem to enable enterprise-oriented Decision Intelligence, but no previous work explicitly identifies the necessary patterns for implementing governance-embedded Decision Intelligence by integrating and optimizing ServiceNow functionality and agent-based systems. The absence of such an unbundling of these systems within a single integrated platform is a major barrier to enterprise-ready Decision Intelligence.

There is substantial interest within a large Financial Services company in an integrated set of autonomous workflows that span departments, regions, products, and levels of criticality. Such workflows are both reliable and

safe, freeing personnel from routine tasks, while allowing the company to weather unforeseen events, disruptions, and crises. A governance-embedded RAG-Agent architecture enabling these Large Language Models (LLMs) is an area of increasing scientific and practical interest, with substantial implications for automating decision-making in high-impact and high-variance environments. ServiceNow is well-placed to serve as the orchestration fabric for these agents. A governance-enabled Data-Informed Self-Learning Risk-Embedded Agentive System for enterprise Digital Decision Intelligence is required to realize these ambitions.

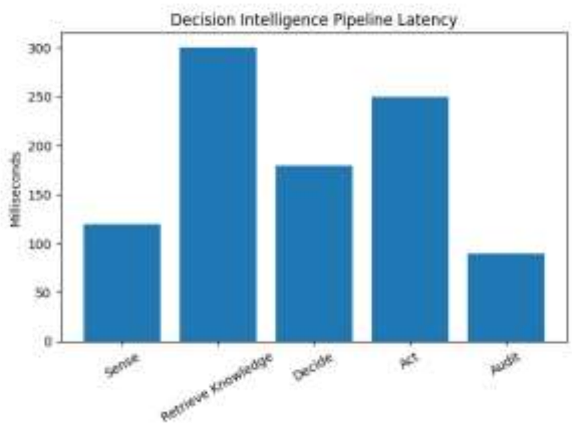

Equation 1: Risk–Agility–Governance Trade-off Model

1.1 Risk Exposure Index (REI)

$$\mathrm{REI}_t = \sum_{j=1}^{m} w_j \cdot \mathrm{impact}_j \cdot \mathrm{likelihood}_j$$

Used at **every pipeline checkpoint** (as described in Sections 6–7).

1.2 Governance-Aware Utility Function

$$U(a_t) = \alpha \cdot \mathrm{Value}(a_t) - \beta \cdot \mathrm{REI}_t - \gamma \cdot \mathrm{ComplianceCost}(a_t)$$

The agent selects:

$$a_t^* = \arg\max_{a \in \mathcal{A}_{\text{safe}}} U(a)$$

2.1. Research Context and Significance

The architecture has been implemented and is currently being tested in a multinational company operating in the fast-moving consumer goods industry. It is designed to govern a family of consolidated ServiceNow workflows related to incident-response processes. The successful execution of these workflows helps protect the organization’s reputation and indirectly safeguards revenue. On the other hand, material missteps in any of these workflows may incur financial penalties, harm brand equity, or attract unwelcome media attention. As such, governance-embedded operational decision intelligence in this domain not only supports standard operating procedures that seek to minimize the likelihood of severe incidents but also helps ensure that decisions made during crisis situations are adequately overseen. This balances the demand for speedy responses with appropriate checks and safeguards.

The work contributes to ServiceNow research by investigating how this platform can facilitate governance-embedded enterprise decision intelligence. Many automated decision-support systems focus on specific tasks and manage their own data, knowledge retrieval, learning, and oversight. As such, they have limited integration with the rest of the organization. ServiceNow offers a rich API surface enabling developers to create, read, and update all kinds of records in addition to orchestrating processes. Therefore, it can underpin operations that extend end-to-end across disparate domains, connecting silos either directly or through more specialized solutions. Moreover, the setup identifies opportunities for greater automation through decision points in existing ServiceNow workflows. For these points, suitable action patterns are catalogued. When these autonomous decisions are possible, a fallback path specifies an alternative strategy to follow. Finally, workflows that once required human intervention or an active approval response can now incorporate suitable safety envelopes.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

3. Conceptual Foundations of Decision Intelligence

Decision intelligence is defined and distinguished from artificial intelligence, machine learning, and decision support. A framework and terminology are introduced to position decision intelligence as a vital enabler of enterprise governance-embedded, risk-aware, autonomous service operations. Governance is further explored as the effective and efficient management of decision-making to achieve corporate objectives and a competitive edge aligned with the enterprise's mission, vision, and values while safeguarding stakeholders' interests.

Decision intelligence—integrating the digitally augmented observation–interpretation–action cycle with enterprise governance to achieve corporate objectives and a competitive edge while safeguarding stakeholder interests—enables enterprises to automate decision-making and service delivery in the digital economy. Digitalisation has catalysed a multitude of technologies that can materially reduce time-to-decision and time-to-action for an enterprise. Nevertheless, the combined and integrated use of these technologies has been limited, as has their alignment to bottom-line business objectives. Decision intelligence seeks to address this need, serving as an organisational framework for the informed end-to-end automation of humans-engaged observation–interpretation-action cycles by embedding AI-enabled, risk-aware, autonomous decision-making into enterprise systems. Enterprises are consequently able to offer decision-making and service delivery—preferably both automated and autonomous—at an increasing rate that is consistent with their mission, vision, and values while also safeguarding the interests of all relevant stakeholders.

Fig 2: Beyond AI: A Decision Intelligence Framework for Governance-Embedded and Risk-Aware Autonomous Service Operations

3.1. Governance in Decision-Making

Governance ensures that decisions are made by the right people, in the right way, at the right time, based on the right information, and with due regard for the risks involved. For enterprise decision intelligence to function correctly, all delegated decision rights must be enforced, accountability must be ensured by defining who must be consulted or informed when decisions are made, and all decision-making roles individuals occupy in a professional capacity must be connected to the enterprise's governance framework.

Governing Decision-Making explains conditions, roles and responsibilities, and decision rights. Decision-making roles must also be clearly outlined with respect to all other governance, risk, and compliance mechanisms that help ensure that decisions are aligned with business objectives, conform to applicable regulatory and legal requirements, and enable appropriate identification and management of

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

the related risks. A number of enterprise policies require human oversight and support of decisions at various stages, which are made clear by the Monitoring Roles in ServiceNow decision table.

4. Architecture Overview

Each direction conceptualizes decision intelligence as a pipeline traversing various data and knowledge sources to derive actionable recommendations. Recent enterprise architectures advocate a governance-enabling environment that orchestrates business activities across application silos. Such a governance-embedded architecture relies on appropriate domain knowledge to determine the rationale behind business goals and assess the impact of any deviation from established policies. ServiceNow, a leading platform for enterprise service management, consists of diverse process automation capabilities—primarily focused on incident, change, and service management—that can govern and manage activities across different business functions even when ServiceNow is not the destination system of record. However, the native product offering lacks the breadth of coverage and depth of intelligence required to enable autonomous decision-making. An integration-oriented approach that leverages the platform's decision points, externalizes decision-critical information, and creates a safety envelope around automated activities would supplement the role of ServiceNow as an orchestration fabric for enterprise operations.
The architecture consists of three planes—control, knowledge, and data—across which the specific components interact. Each plane supports a well-defined role that can be leveraged by the system design. The control plane orchestrates end-to-end operations, enables communication between various components, and defines the distinct roles of the involved roles; the knowledge plane stores domain-specific knowledge and governs the decisions being taken; and the data plane consolidates the data required for inference, action selection, and action effect assessments based on the historical activity audit trail. Building enterprise decision intelligence that supports intelligent self-service, enables action without direct human intervention, safeguards against potential risk, and exploits data available across the enterprise data landscape acts as the main criterion for a successful design.

4.1. System Goals and Constraints
The central aim of the proposed ecosystem is to enable autonomous ServiceNow workflows without sacrificing the foundations of decision-making: proper governance embedded in policies, rules, and initiatives that ensure compliance with established norms. The primary characteristics that govern the design of these enterprise processes are their reliability, security, adherence to regulatory, legal, and quality constraints, and sufficient capacity to handle the expected workload.
To be deemed trustworthy, the architectural components must operate as expected without any catastrophic issues: they must provide the expected capabilities across functional areas, deliver the right information in a timely manner, and share information securely with authorized actors only while mitigating the risk of unauthorized access. Furthermore, decision-making must comply with relevant organizational policies and regulations, ensuring that it controls sensitive data, shielding its usage from unwarranted exposure and data hoarding, and allowing the rationales behind decisions to be accessible for cross-examination by all involved stakeholders. Moreover, all external operations must be protected from malicious actions. Finally, compliance with the service-level agreements defined for supported services must be guaranteed.

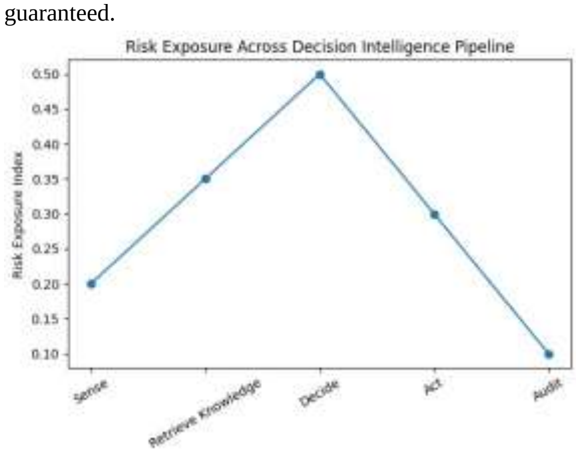

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

Equation 2: Knowledge Retrieval and Vectorization (RAG Layer)

Governance-embedded Retrieval-Augmented Generation.

2.1 Vector Embedding Function

$$\mathbf{v}_i = f_{\text{embed}}(x_i), \quad \mathbf{v}_i \in \mathbb{R}^n$$

Where:

- x_i: document, incident, or policy text
- f_{embed}: embedding model (LLM encoder)

2.2 Policy-Constrained Retrieval

$$\mathcal{K}_t = \underset{\mathcal{K} \subseteq \mathcal{D}}{\operatorname{argmax}} \sum_{k \in \mathcal{K}} \text{sim}\,(\mathbf{q}_t, \mathbf{v}_k) \quad \text{s.t. } k \in \mathcal{P}_{\text{allowed}}$$

This enforces **policy-bounded retrieval**, a key contribution of your architecture.

5. ServiceNow as an Orchestration Fabric

While ServiceNow possesses multiple characteristics that enable decision intelligence goals around RAG and agent-based architecture, it also exhibits constraints, particularly with respect to standalone autonomous decision-making. Nevertheless, sufficient facilities exist within its low-code framework to enable rapid decision deployment, bolstered by a wealth of existing infrastructure and domain knowledge with multiple data integration points and API surfaces for access to data silos, supplemented with third-party plugins for advanced capabilities.
ServiceNow functions as the decision guidance platform facilitating RAG and agent-based architecture without introducing autonomous decision-making. Data integration points exist across multiple enterprise solutions supported via API integration for multi-modal data access. Knowledge around historical incidents, service change, service fulfillment, and associated risk attributes is either natively available within ServiceNow or is accessed via API integration. Hence, SafeCloud's governance-embedded Decision Intelligence Lever is implemented for complete Decision Intelligence within the enterprise SecureCloud™. ServiceNow internal capabilities and low-code platform enable the rapid creation of SharePoint-like sites, with multiple embedded decision orchestration workflow patterns available out-of-the-box. Each pattern maps out an end-to-end service interaction flow; however, the decision function is limited to a sub-process within the overall workflow. When ServiceNow functions as the overarching governance fabric, additional nodes enable and protect against risks and repeatability of decisions through multiple fallback options and safety envelopes.

5.1. Data Silos, Integration Points, and API Surface
A comprehensive inventory of data silos highlights various dissociated data repositories across the enterprise. Governance surface mappings of the distinct data products indicate areas where integration is deemed highly desirable given its associated business value. These range from relatively isolated data silos with only an IETF-approved data sharing protocol to ServiceNow integration points where both pull and push services expose contract conditions covering the standards of accessibility, …, and other M2M conditions,… beyond the well-understood MAPI security and non-repudiation. Other DataSilk … integration points provide natural extensions in their data availability, …, or being indirectly or directly connected to population-critical events: e.g., Stocke > Daily Management IR and MsOffice reports, Product and people quality rating APIs, ... etc.
ServiceNow's substantial API surface is an additional attribute on ServiceNow's Capability Surface warranting detailed evaluation investments. The annotation and analysis of prior pattern-based deployments point to Data Governance Books equipping the available APIs. Spherical mapping of IC-PM and IC-Argos data Products evaluating their informational quality index and production stability of all DataChemOps contribute towards driving operational research into automating its usage during policy compliance checks between customer and supplier organizations in an enterprise context.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

Fig 3: Autonomous Data-Governance Synthesis: Spherical Mapping and API Orchestration for Resolving Enterprise Silos through Policy-Driven DataChemOps

5.2. Autonomous Workflow Design Patterns within ServiceNow

Using these qualities, it is possible to identify ServiceNow's out-of-the-box automation patterns and operations that can function without human involvement. Importantly, these patterns will support automated workflows with internal decisions that lead to human intervention only if considered necessary by the patterns. Such patterns must also ensure safety, including human fallback options in the event of adverse scenarios. Thus, it is essential to define these operational patterns in detail while mapping out the corresponding decision points and fallback behavior.

Patterns considered for inclusion in the autonomous decision part of the pipeline during this assessment encompass those decision operations where an intelligent execution segment can direct other ServiceNow functions to assist external actors, and an internal decision must be implemented based on success/failure. These patterns also contain the potential for future operators to reduce resource expenditure further by automating low-impact fields of Day 2 and OMG operations. In addition, they cover areas characterized by the presence of a human operator whose responsibility is merely to approve actions that the pattern has determined to be within their delegation of authority. These cover different areas ranging from Unified Monitoring into Day 2 and OMG activities, Operations/Facilities, and Hybrid Cloud Management.

6. Decision Intelligence Pipeline

The complete flow of the Decision Intelligence ecosystem—from sensing through data processing, analysis, decision-making, action initiation, and execution—embodies the end-to-end process from input to outcome. Explicit governance embeds institutional policy within the pipeline at multiple stages, thus ensuring that the right decisions are made and executed by the right agents, in accordance with enterprise policy, and fully traceable to a definitive source of truth. Governance mechanisms safeguard the ecosystem before, during, and after execution, ensuring that the right checks and balances are in place.

Decision Intelligence operates as a pipeline or pipeline of pipelines, encompassing processes that traverse knowledge retrieval, analysis, insight generation, and action initiation. These processes include fundamental operations for nearly all classes of decision-making—from Alerting, Acknowledging, Assessing, Resolving, to After-Action Review. At each stage of this data-to-decision pipeline (DP2D), governance checkpoints verify risk exposure against defined appetite levels, such as assessing the potential impact of a cyber incident before instructing response teams to initiate remediation actions or modulation levels. During execution, further safeguard mechanisms can remain in place, re-routing operations to avoid undesirable (but not fully prohibited) actions when a safety envelope is deemed to be breached. Finally, disparities between service performance and governance risk profiles may trigger a closed-loop feedback process, steering the system and its behaviours towards the desired modes of operation and performance.

Equation 3: Decision Function with Governance Safeguards

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

3.1 Decision Policy

$$a_t = \pi(s_t \mid \theta)$$

Where:

- a_t: selected ServiceNow action
- π: agent policy (RAG-agent)
- θ: learned parameters

3.2 Safety Envelope Constraint (SLAAAP)

SLAAAP (Sensitive, Legal, Accounting, Approval, Policy).

$$a_t \in \mathcal{A}_{\text{safe}} \Leftrightarrow \bigwedge_{i \in \text{SLAAAP}} C_i(a_t, s_t) = \text{true}$$

If violated:

$$a_t \to a_t^{\text{fallback}} \quad \text{(human-in-the-loop or escalation)}$$

6.1. Knowledge Retrieval and Vectorization

To embed knowledge-driven models into the deployed architecture, the subsequent processing stage ingests data from the relevant knowledge and data silos, generating the input artifacts necessary to enable the triggering of autonomously defined ServiceNow workshops. The triggering represents the 'Push' step of the Pipeline depicted in the architecture overview. Based on the utilized knowledge brokers, the appropriate models are engaged to push the corresponding input to the workshop and start the ServiceNow activity.

In the context of governance-embedded Decision Intelligence, the following conditions apply: (1) for a Knowledge Broker that is a Question Answering Model, any of its definition or usage attributes can be ingested—regardless of whether they are deemed secure, private, or related to sensitive topics—by the ServiceNow Governed Data Knowledge data source. The only restriction is that the content needs to be used for training the making decisions Workshop class Age of the Agent; (2) for a Knowledge Broker that applies the five W's and an H rule (Who, What, Where, When, Why, and How) during its operation, a specific rule needs to be defined in order to map the input stemming from the Governed Data Knowledge data source to the corresponding Who, What, Where, When, Why, and How; (3) for a Knowledge Broker that is a Visual Tool Creator, any user-generated content that is not marked as private or secure can be ingested by the ServiceNow Governed Visual Knowledge data source; (4) for a Knowledge Broker that is a Visual Representation Model, any of the visual definitions can be used as input for firing the Visual Activities Workshop class pattern when training the making decisions Workshop class Age of the Agent.

7. Governance-Embedded Safeguards

Safeguards embed policy and oversight mechanisms within the decision intelligence pipeline, ensuring compliance with governance frameworks while enabling responsive decisions. Naturally, the goal of enabling governance is to enhance risk control, with the corresponding trade-offs being a higher management burden and possibly reduced agility. Thus, the primary differentiation of the architecture is the integration of safeguards – without these elements, the underlying logic could be replicated with near-identical results.

The first safeguard focuses on explicitly defined policies that control behaviour. No pipeline component – including cognitive modules, underlying LLMs, or V&F Data – is exempt from policy definition and systematic enforcement. Mediation rules, expressed in natural language, are automatically transformed into appropriate actions at different points within the tackle. It is the role of policies to reduce the probability of significant incidents by allowing only a regulated state space and set of actions.

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

Fig 4: Architectural Safeguards in Decision Intelligence: Systematic Policy Enforcement and Regulated State Spaces for Autonomous Governance

7.1. Policy Definition and Enforcement

Lateralization of command and control not only accentuates the imperative of risk governance as a prerequisite for decision-making, it also inherently separates decision-making from policy formulation. Within the context of an operating environment like ServiceNow, this is summarily addressed within ServiceNow Policy and Compliance Management. Here rules are formulated as part of a broader GRC narrative and mapped to actions that rely on ServiceNow workflows. The content lifecycle then enables the automated enforcement of these rules across ServiceNow and, where a boundary or interface to the ServiceNow operating fabric is defined, also beyond ServiceNow.

Implicit in this separation of roles is that the mapping of the policy is done by the governing entity, however this is not prescriptive in nature: during corrective cycles it is equally feasible for the detection layer to formulate the rules based on Non-Conformity Reports (NCRs), Demand, and Request fulfillment which are also recorded within the Policy and Compliance Management module.

Stage	Latency (ms)	Risk Exposure Index
Sense	120	0.20
Retrieve Knowledge	300	0.35
Decide	180	0.50
Act	250	0.30
Audit	90	0.10

7.2. Auditing, Transparency, and Traceability

Comprehensive audit trails provide a detailed record of events for verification and prevention of malicious behavior. Every request, response, decision-making action, provenance information, and explanation of rationale must be logged in an immutable storage. The audit trails must satisfy the principles of explainability defined in the subsequent section. Representing the various decision-making calculation models as features in a decision-making tree enables users impacted by the decision to make sense of how the AI engine reaches a specific decision. Audit trails facilitating transparency empower the users to hold the DI workflow accountable for its decisions. Critical decision paths sensitive to potential regulatory compliance must provide detailed explanations for every decision. The data must be stored in an explainable and searchable format to facilitate compliance reporting.

The auditors should be able to interrogate the audit trail to find out the route for any category A or category B decision as defined in the service policy engine. All AI-enabled decisions must be linkable back to the request that initiated the decision. Traceability is the ability to follow a decision backward beyond the direct trigger. For example, if an AI-enabled decision is to suspend a user account in response to suspected malicious behavior, it should be possible to link the decision back to the persistence of the user and the associated category D Security & Compliance policy that has driven the status of the user to suspected-malicious-behavior.

8. Conclusion

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

Decision agents enhance decision-making speed and provide a competitive edge by performing lower-risk and non-strategic functions independently. Although the pressure to automate everything is strong, failure or wrong expectations when workflows go completely autonomous can be detrimental. Unexpectedness in automatic decision-making or skill application can be harmful to both the enterprise's financial bottom line and brand value. Consideration, capacity, and budget constraints can minimize the underlying risks, allowing parts of a workflow to run autonomously but still under the supervision of human stakeholders. Decision agents can be employed to autonomously steer subprocesses of such workflows within a SLAAAP (Sensitive, Legal, Accounting, Approval, and Policy-sensitive) envelope or on customized paths at predefined decision points. ServiceNow fulfills the role of an intelligent workflow automation platform. The combination of enterprise governance, risk, value, and regulatory compliance with automation in ServiceNow therefore sets the stage for truly intelligent enterprise operations.

This contribution put forth a high-level design of a governance-embedded intelligent ServiceNow architecture, leveraging reinforcement learning from human feedback and retrieval-augmented generation principles. By fulfilling critical requirements such as embedding regulation and oversight into each stage of the autonomous decision-making pipeline, the proposed architecture ensures appropriate levels of enterprise, data, privacy, and security governance. Through private information retrieval, it protects sensitive data and aspects of business know-how during external consultations. Future work will build the individual pieces, demonstrate their working, validate their usefulness, and put the architecture into practice.

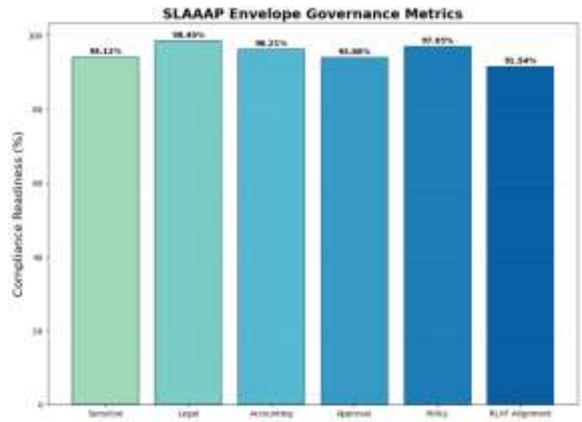

Fig 5: SLAAAP Envelope Governance Metrics

8.1. Final Thoughts and Future Directions

Formalization of Decision Intelligence highlights a growing doctrine within large organizations. Current tech and trends, amplified by ChatGPT and friends, have created thousands of applications and implementations of less than reliable solutions. Use case prioritization and causation validation now appear mandatory to avoid rushing to production and falling prey to the transitoriness that underminds hype-fuelled. The here proposal embeds intelligence into a leading system of record decision intelligence, while still remaining capable of bridging the Enterprise with other silos of data. Should such solutions ever reach production scale the requirement and or command of designed Governance controls along the information life-cycle becomes priority.

Living and Change in the Open Source World. Governance should never stop protecting what matters to a Society nor should remain static. ChatGPT and friends have multiplied the usage of languages and engines. These available at low or no cost can now perform task that used to be reserved to experts. From the creation of hobby or vacation articles to the generation of complex and coherent page data has become cheaper than it was 20 years ago. Still the quality that drives readers to pay for purchasing not only Creation but the selection and validation process still stands. Change is widely accepted, embraced and lived within the Open Source society. It yet becomes evident that Governance controlling the usage of such technologies should not aim

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

to stop change but rather to control side effects and negative output, thus generating a happiness enabler rather than burden.

9. References

1. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N. S., Chen, A., Creel, K., Davis, J. Q., Demszky, D., ... Liang, P. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258.
2. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 610–623). Association for Computing Machinery.
3. Kiela, D., Bartolo, M., Nie, Y., Kaushik, D., Geiger, A., Wu, Z., Vidgen, B., Prasad, G., Singh, A., Ringshia, P., Ma, Z., Thrush, T., Riedel, S., Waseem, Z., Stenetorp, P., Jia, R., Bansal, M., Potts, C., & Williams, A. (2021). Dynabench: Rethinking benchmarking in NLP. In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (pp. 4110–4124). Association for Computational Linguistics.
4. Valiki, D., & Segireddy, A. R. (2023). Deep Learning Architectures Deployed on Cloud Platforms for Dynamic Financial Risk Evaluation and Market Prediction. American International Journal of Computer Science and Technology, 5(5), 12-24.
5. Askell, A., Bai, Y., Chen, A., Drain, D., Ganguli, D., Henighan, T., Jones, A., Joseph, N., Mann, B., DasSarma, N., Elhage, N., Hatfield-Dodds, Z., Hernandez, D., Kernion, J., Ndousse, K., Olsson, C., Amodei, D., Brown, T., Clark, J., ... Kaplan, J. (2021). A general language assistant as a laboratory for alignment. arXiv preprint arXiv:2112.00861.
6. Borgeaud, S., Mensch, A., Hoffmann, J., Cai, T., Rutherford, E., Millican, K., Van Den Driessche, G., Lespiau, J.-B., Damoc, B., Clark, A., De Las Casas, D., Guy, A., Menick, J., Ring, R., Hennigan, T., Huang, Z., Maggiore, L., Jones, C., Cassirer, A., ... Sifre, L. (2022). Improving language models by retrieving from trillions of tokens. In Proceedings of the 39th International Conference on Machine Learning (pp. 2206–2240). PMLR.
7. Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. In Advances in Neural Information Processing Systems, 35, 27730–27744.
8. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q. V., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. In Advances in Neural Information Processing Systems, 35, 24824–24837.
9. Inala, R. (2020). Building Foundational Data Products for Financial Services: A MDM-Based Approach to Customer, and Product Data Integration. Universal Journal of Finance and Economics, 1(1), 1-18.
10. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). ReAct: Synergizing reasoning and acting in language models. arXiv preprint arXiv:2210.03629.
11. Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., Olah, C., Hernandez, D., Drain, D., Ganguli, D., Li, D., Tran-Johnson, E., Perez, E., Kerr, J., Mueller, J., ... Amodei, D. (2022). Constitutional AI: Harmlessness from AI feedback. arXiv preprint arXiv:2212.08073.
12. Liang, P., Bommasani, R., Lee, T., Tsipras, D., Soylu, D., Yasunaga, M., Zhang, Y., Narayanan, D., Wu, Y., Kumar, A., Newman, B., Yuan, B., Yan, B., Zhang, C., Cosgrove, C., Manning, C. D., Ré, C., Acosta-Navas, D., Hudson, D. A., ... Hashimoto, T. (2022).

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

Holistic evaluation of language models. arXiv preprint arXiv:2211.09110.

13. Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. arXiv preprint arXiv:2303.12712.
14. Inala, R. (2023). Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective. International Journal of Finance (IJFIN)-ABDC Journal Quality List, 36(6), 579-606.
15. Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W. X., Wei, Z., & Wen, J.-R. (2023). A survey on large language model based autonomous agents. arXiv preprint arXiv:2308.11432.
16. Schick, T., Dwivedi-Yu, J., Dessi, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). Toolformer: Language models can teach themselves to use tools. In Advances in Neural Information Processing Systems, 36.
17. Asai, A., Wu, Z., Wang, Y., Sil, A., & Hajishirzi, H. (2023). Self-RAG: Learning to retrieve, generate, and critique through self-reflection. arXiv preprint arXiv:2310.11511.
18. Shinn, N., Cassano, F., Berman, E., Gopinath, A., Narasimhan, K., & Yao, S. (2023). Reflexion: Language agents with verbal reinforcement learning. In Advances in Neural Information Processing Systems, 36.
19. Park, J. S., O'Brien, J. C., Cai, C., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behavior. In Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology, 1–22. Association for Computing Machinery.
20. Aitha, A. R. (2021). Optimizing Data Warehousing for Large Scale Policy Management Using Advanced ETL Frameworks.
21. Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T. L., Cao, Y., & Narasimhan, K. (2023). Tree of thoughts: Deliberate problem solving with large language models. In Advances in Neural Information Processing Systems, 36.
22. Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., Awadallah, A. H., White, R. W., Burger, D., & Wang, C. (2023). AutoGen: Enabling next-gen LLM applications via multi-agent conversation. arXiv preprint arXiv:2308.08155.
23. Li, G., Hammoud, H. A. K., Itani, H., Khizbullin, D., & Ghanem, B. (2023). CAMEL: Communicative agents for “mind” exploration of large scale language model society. arXiv preprint arXiv:2303.17760.
24. Patil, S. G., Zhang, T., Wang, X., & Gonzalez, J. E. (2023). Gorilla: Large language model connected with massive APIs. arXiv preprint arXiv:2305.15334.
25. Qin, Y., Liang, S., Ye, Y., Zhu, K., Yan, L., Lu, Y., Lin, Y., Cong, X., Tang, X., Qian, B., Zhao, S., Hong, L., Tian, R., Xie, R., Zhou, J., Gerstein, M., Li, D., Liu, Z., & Sun, M. (2023). ToolLLM: Facilitating large language models to master 16000+ real-world APIs. arXiv preprint arXiv:2307.16789.
26. Liu, X., Yu, H., Zhang, H., Xu, Y., Lei, X., Lai, H., Gu, Y., Ding, H., Men, K., Yang, K., Zhang, S., Deng, X., Zeng, A., Du, Z., Zhang, C., Shen, S., Zhang, T., Su, Y., Sun, H., ... Tang, J. (2023). AgentBench: Evaluating LLMs as agents. arXiv preprint arXiv:2308.03688.
27. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Canton Ferrer, C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., ... Scialom, T. (2023). Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288.
28. Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., & Liang, P. (2023). Lost in the middle: How language models use long contexts. Transactions of the Association for Computational Linguistics, 11, 1576–1592.
29. Nagabhyru, K. C., & Engineer, S. D. (2023). Unifying Data Engineering and Machine Learning Pipelines: An

ISSN: 3067-4190
JOURNAL OF ARTIFICIAL INTELLIGENCE AND BIG DATA DISCIPLINES
Volume: 02 Issue: 02
RECEIVED: APRIL 23 REVISED: MAY 04 ACCEPTED: MAY 25 PUBLISHED: JUNE 20

Enterprise Roadmap to Automated Model Deployment.

30. Naveed, H., Khan, A. U., Qiu, S., Saqib, M., Anwar, S., Usman, M., Akhtar, N., Barnes, N., & Mian, A. (2023). A comprehensive overview of large language models. arXiv preprint arXiv:2307.06435.