




Governed Agentic AI for ITSM and HRSD Automation

Dimitris Plexousakis
Asst Professor

Abstract


Governance-native agentic AI provides a promising architecture paradigm for enterprise automation needs. The concept is placed into context through exploration of the role of governance in AI development, supported by a framework for designing and deploying such systems. The architecture is then exemplified through definition of requirements and an RAG pipeline development approach for IT Service Management (ITSM) and Human Resource Service Delivery (HRSD) operations, with particular focus on the practical establishment of RAG systems for knowledge-rich conversational contexts. The discussion highlights the novelty of the initiative and its support of implementation of enterprise-scale nudges.

Artificial Intelligence (AI) is now a key enabler of corporate automation initiatives. Such systems perform a myriad of tasks, particularly in the Information Technology (IT) and Human Resources (HR) functions of enterprises, and are widely presented with autonomous agent labels. However, while the AI product offerings of the major technology companies have grown in sophistication, the accompanying Agentic Operating Systems (AOS) technologies are far less developed. Hence, few AI systems are actually agentic in functioning, and even formal agents sense decisions and actions assigned within the bounds of a specific conversation context.

Keywords: Governance-Native Agentic AI, Enterprise Automation Architecture, AI Governance Frameworks, Agentic Operating Systems (AOS), Retrieval-Augmented Generation (RAG), Knowledge-Rich Conversational Systems, IT Service Management Automation (ITSM), Human Resource Service Delivery (HRSD), Enterprise-Scale Nudging, Governance-by-Design AI, Agentic AI Deployment Frameworks, RAG Pipeline Engineering, Conversational Knowledge Orchestration, Autonomous Enterprise Agents, Decision–Action Loop Design, Responsible Agentic AI, Corporate Automation Enablement, Context-Aware AI Agents, Operationalizing Governance in AI, Enterprise Conversational Platforms.

1. Introduction

This paper outlines a governance-native agentic AI architecture capable of synthesizing multiple retrieval-augmented generation (RAG) pipelines for complex enterprise tasks.



Enterprise IT service management (ITSM) and human resources service delivery (HRSD) are considered exemplary use cases. Enterprise ITSM relies on numerous sources of interactive support data, including access management and assurance documentation, and other operational knowledge. HRSD spans a range of specialist areas such as recruitment, employee relations, and learning and development. Multiple natural-language-based RAG pipelines are therefore required, and their overall orchestration must be governed, guided, and verified holistically, sustainably, and resiliently over time through a combination of orchestrator roles, decision policies, and appropriate interactions with humans and other agents.

Driven by the firm conviction that AI can play a decisive role in the future of enterprise ITSM and HRSD and, indeed, any domain requiring complex synthesis of interactive support or service knowledge from multiple structured and unstructured sources, the capability to architect into AI systems such native governance and agentic qualities, in a practical context, is deeply compelling. Enterprise ITSM and HRSD are particularly appealing because of their expansive and fully supported knowledge resources and precedent case studies, both of which are vital for underpinning the design of advanced RAG pipelines and their orchestration.

1.1. Overview of Key Concepts and Objectives

A perspective on platform-native agent-based big-brain development of independently intelligible, derivatively discoverable, dynamically constituted and deployed, capability-based, domain-contributory, distortion-tolerant, Lexicon-as-a-Service underpinnings for risk-balanced, gated and supervised-orchestrated RAG pipelines for the trusted realisation of return-risk-resilient, credibility-preserving, viable-to-dependable, human-ready generative intelligence: opening up information technology supply chains and service functions, delivering risk-reward-

nexus-aligned automation of fully operational ceiling-defining ITSM and HRSD processes.





Fig 1: Agentic Lexicon-as-a-Service (LaaS): A Risk-Resilient Framework for Supervised RAG Pipelines in Enterprise ITSM and HRSD Orchestration.

Guiding principles influence and shape the design, implementation and governance of any technology-based solution; these principles must be intrinsically understood, realised, facilitated, mediated or enforced, as appropriate, by any unarguably genuine or native realisation




of such a technology-based solution — and, in a context where the applied development of large language models is presenting both opportunities and challenges for the Information Technology Management function of enterprises, the informed application of this requirement has motivated the development of a conceptual perspective on a specific platform-native capability proposition that is, at its core, an instantiation of risk-reward-nexus-aligned, credibility-preserving, agentic artificial intelligence that is designed to define and deliver solutions with syntactically and semantically intelligent, risk-balanced, gated and supervised-orchestrated Generative RAG implementations for the trusted exploitation of an underlying information content source bank.

2. Background and Motivation

An argument for AI governance is mounted, focussing on three aspects: contextualization of governance within responsible AI, characterization of the currently constraining role attributed to governance in AI development, and evaluation of the risk posed by autonomous AI overstep of governance-imposed boundaries. Governance matters; regardless of the openness or closedness of the model development process, without the necessary level of pragmatism, the risks posed by the adoption of AI by business, such as consequent biases, legal and financial risks, inappropriate corporate decisions, fake news, etc., will prevail over its benefits. Indeed, far from being a comfortable shield, governance remains a critically constraining factor that entails risk for all those whose primary aim is to deliver agentic AI, who must situate themselves within a broader governance structure granting oversight and control of the acceptability of not only the development process but also of the operations of the AI agents once deployed. But even such control does not totally eliminate risk: the correctly framed problem still has the potential to generate signals so compelling that the agentic AIs will want to act on them and be tempted to overstep their boundaries.

The primary argument focuses on the internal axis of the responsible AI triangle, to illustrate how well-designed pipelines can ensure that agentic AIs remain an expression of the expectations of the designers, the business and society. As agentic AI systems gain autonomy over decision-making, the combination of a robust design of the problem being treated with well-defined pipelines still guarantees a control mechanism, through the decision policies collectively defining the agentic behaviour of the deployed AIs, that enables sufficient supervision, timely intervention and forced escalation when crossing red lines.



Equation 1) RAG as a probabilistic mixture: generation grounded on retrieved evidence

Let:

- User query = q
- Retrieved evidence chunks = $c_1, c_2, \dots, c_k$
- Output answer tokens = $y = (y_1, \dots, y_T)$

Goal: model the probability of an answer given the query, *using retrieved context*:

$$P(y | q) = \sum_C P(y | q, C) P(C | q)$$

where C is a set of retrieved chunks.


Step-by-step intuition

1. We don't want a pure generator $P(y | q)$ because it can hallucinate.
2. RAG introduces a latent variable C ("the evidence we condition on").
3. Apply the law of total probability over C :

$$P(y | q) = \sum_C P(y, C | q) = \sum_C P(y | q, C) P(C | q)$$

4. In practice, approximate by top- k retrieval results:


$$P(y | q) \approx \sum_{i=1}^k P(y | q, c_i) P(c_i | q)$$



2.1. Contextualizing Governance in AI Development

Governance-Native Agentic AI Architecture for Enterprise ITSM and HRSD Automation via Orchestrated Retrieval-Augmented Generation Pipelines. The AI boom presents unprecedented opportunities for automated enterprise operations, yet it has become clear that broad talent-skilling and talent-outreach operations generate only limited business-relevant value without strict output governance. An architecture is proposed for governance-native agentic AI, designed specifically for enterprise IT service management (ITSM) and human resource services delivery (HRSD) functions. Automation is achieved through orchestrated retrieval-augmented-generation (RAG) pipelines, each tailored to a specific production or consumption context.

Contextualization of Governance in AI Development. The broad availability of AI capabilities pretrained on public web data, combined with the corresponding boom in the development and deployment of AI-based government services for talent skilling, outreach, and support, presents opportunities for organisations to automate key enterprise functions such as human resource services delivery (HRSD) and IT service delivery management (ITSM). However, the ungoverned and uncontrolled use of pretrained models has resulted in limitations



related to privacy, confidentiality, safety, trust, and reliability, the failure to observe public administration norms and the replicability of the public-private model.

The architecture proposed in this work is focused on the context of such unregulated public AI consumption and is subsequently adapted to enterprise ITSM functions and other agents requiring validated responses. Such functionality goes beyond simple fact-checks, enabling verification and validation against trusted internal and enterprise private knowledge bases, deployment in knowledge domain systems, and dynamic retraining and tuning of large language models (LLMs).


3. Conceptual Framework: Governance-Native Agentic AI

A governance-native agentic AI architecture is proposed with two primary features. First, the architecture is designed from start to finish with governance-as-a-principle in mind: the support of ethical and responsible AI development as a primary contextual consideration. The second feature arises by subdividing the agentic AI systems into constituent components, roles, and design considerations, and then conducting the architecture design by addressing each at every relevant layer of the technology stack. The resulting architecture has two application domains, enterprise IT service management (ITSM) and human resources service delivery (HRSD), and combines multiple retrieval-augmented generation (RAG) pipelines in an orchestration layer.

Two primary groups of service requests are identified in the context of enterprise IT service management (ITSM): those that require human decision-making and those that are suitable for a decision rule-based system. Within ITSM, the context and requirements of a Governance, Risk and Compliance (GRC)/IT-general control-related Request-For-Information (RFI) related recurrent request type is presented as an example. The context and requirements for the human resources service delivery (HRSD) domain involve the orchestration of multiple general-purpose RAG pipelines with an emphasis on quality, data monitoring, and control to address relational and sensitive requests, while also supporting transactional service request automation.

3.1. Principles of Governance in AI Systems

Techniques for AI-enhanced governance in organizations and software development generally can also be applied to the subfield of AI creation itself, especially for AI-systems early



stages not focused on governance e.g., business management and direction. Early AIs are not in charge of handling open-ended consequences of misleading decisions by themselves (e.g., generating unintended consequences). Just as a governance entity needs a board, an AI system that requires more agency than simple answering needs an agentic RAG pipeline for guidance.

The principles of AI system governance do not depend on those for Governance-Native Agentic AI. The first principle states that the usage [activation] of AI technology should remain bounded by established limitations on its application. This applies to training and testing of generative LLMs but not to generation-completion systems like Codex and ChatGPT. Governance-nature AI technology is constantly being improved ideally without supervision of multiple teams and other decision bodies. The second principle states that AI enhancements should only be used for value-additional tasks where generation-completion capacity cannot harmoniously satisfy the requirements without additional AI activation. E.g.: there is neither benefits nor losses when using an AI-enhanced library like Codex inside an AI-enhanced code writing generation-completion application. The third principle states that AI activations must be automatable and standardized where numerous activations are required. In this sense, using DALL-E inside a prompt-guided application aimed at prompt construction, is clearly not satisfying the localization capacity of the AI system since the prompt for DALL-E has zero informative value.





Fig 2: Governing the Genesis: Architecting Agentic RAG Frameworks and Tri-Primal Protocols for Early-Stage AI Development

3.2. Agentic Capabilities and Roles

Various settings, together with the temporal, jurisdictional, and legal frameworks involved, will determine the appropriate balance between the level of autonomy offered to an AI system and the corresponding human degree of oversight. Autonomy must be approached with caution, especially where responsibility for actions or consequences falls outside the algorithm's remit. Such areas include any developments that infringe on the rights or welfare of others; unambiguous breaches of privacy; the formation of beliefs, perceptions, or thoughts that may impact the well-being of the receiver; any event for which laws or regulations impose responsibility on humans; and any action or decision that can cause social backlash or violate cultural norms or context.



Acting within these temporal, jurisdictional, or legal limits, AI can step beyond decision support and, assisted by suitable orchestration, take direct actions on its own. Its agency at this stage resembles that of programmed robots: given constraints and a decision-generation capability, it performs autonomously but incurs no moral blame or legal responsibility because a human offered it the power to act. Unlike programmed robots, AI systems remain bereft of builtin knowledge about how to fulfil a request or act in a context beyond conditions specified beforehand. Yet an orchestrator can configure externally founded, equilibrium-exploiting, exploratory, or learnt pipelines that can act in areas unaddressed at instantiation, enabling within-contract agency akin to regulated discretion. In these cases, the agent acts on behalf of the human and demands the minimum level of supervision warranted.


4. Architecture for Enterprise ITSM and HRSD

While both AI Service Providers and Large Language Models (LLMs) offer innovative technology, they fall short of directly supporting Autonomous and Agentic AI for the Automation of Intelligent IT Service Management (ITSM) and Human Resource Service Delivery (HRSD) Processes. As a result, these enterprises must adapt orchestrated Retrieval-Augmented Generation (RAG) Pipelines to their requirements and limitations within a Governance-Native Architecture Framework. The AI Open Universe serves as the Enterprise ITSM Data Universe, while the AI Private Universe and Enterprise Knowledge Bases act as the Business-Specific Information Sources for the AI Service Provider responsible for ITSM. In the HRSD context, the Enterprise AI Knowledge Base supplies Business-Specific Information and the Data Universe contains Restricted Access natively within the AI Private Universe.

The required multi-agency support for ITSM is organized into three Principal Agent Layers, with multiple Agent Instances at the Enterprise Designated Agent Level. An additional Facilitating Agency overlays the Principal Layer to Support Enterprise AI LLM Data Privacy and Data Security. The resulting Closed Ecosystem and Controlled Environment Limit the Connections with External AI Service Providers. The Agentic Support for HRSD is organized as a Decision-Delegation Matrix based on formal Role, Responsibility, and Accountability Allocation.

4.1. Enterprise ITSM Context and Requirements


IT Service Management (ITSM) provides the conceptual structure and processes to execute and manage IT services requested by employees in their roles as users of internal



systems and applications. This involves managing the delivery of IT services to end-users, employees in enterprise environments. Generic IT services include restore service requests (for mental health issues), employee relocation requests, equipment requests from IT or facilities, missing payroll service requests, etc. These services are dependent on IT apps that prepare for internal controls, manage compliance violations, and help implement company policies.

Governance-Native Agentic AI Architecture for Enterprise ITSM and HRSD Automation via Orchestrated Retrieval-Augmented Generation Pipelines. ITSM services are constrained by the corporate governance structure and organization policies of the enterprise. These usually define the business timelines for service completion, the expertise to investigate a service request, the ability to approve a service fulfillment or completion, etc. Agents cannot operate outside an organization's policy boundaries, which might require explicit approval from another internal system that holds the required authority. Approvals might depend on contextual information, which can form the key components in a prompt to a generation mechanism, retrieved from an app or a database.

ITSM automation can interact with several internal systems, managing and embedding their internal controls and policies in SCM and R&D services, preparing procedural templates and automation tools via systems management, forwarding to the appropriate area the handling of operations-request tickets from subject matter experts, orchestrating apps that require actions from other areas in an enterprise using a service request form like PowerApps, and covering common inquiries based on past elaborations using internal knowledge-apps.



Equation 2) Embeddings + cosine similarity (typical retrieval scoring)

Let:


- Query embedding $e_q \in \mathbb{R}^d$
- Chunk embedding $e_i \in \mathbb{R}^d$

Cosine similarity

$$\cos(e_q, e_i) = \frac{e_q \cdot e_i}{\|e_q\| \|e_i\|}$$

Step-by-step

5. Dot product:



$$e_q \cdot e_i = \sum_{j=1}^d (e_q)_j (e_i)_j$$

6. Norms:

$$\|e_q\| = \sqrt{\sum_{j=1}^d (e_q)_j^2}, \quad \|e_i\| = \sqrt{\sum_{j=1}^d (e_i)_j^2}$$

7. Normalize dot product by magnitudes (so long vectors don't dominate).


Top-k retrieval rule

Retrieve $\arg\max_i \cos(e_q, e_i)$ (pick top k)

4.2. HRSD Context and Requirements

In the HRSD domain, organizations seek to enhance the user experience for external customers and internal staff via streamlined self-service functions. A governance-native agentic AI capability in this area is proposed, addressing both internal and external HRSD activities, such as providing straightforward information on Qualification and Entitlement requirements. In agencies and organizations with broad ITSM responsibilities, the steps required to receive assistance or engage in responsive communication generally engage individual case owners. Moreover, communication with individual providers, when necessary, must provide adequate context. Therefore, an agentic AI system should be designed to engage with external customers and with internal HR providers for a specified area of responsibility. As documented in Section 2, this engages the external service delivery for HRSD and ITSM and supports both internal cloud and temporal oversight delivery.

At a functional level, the overall objective of HRSD systems is to promote autonomy for staff seeking work-related motivation or information assistance that is straightforward and not confidential (e.g., by showing contacts and requirements). In such areas of the business, service owners must ensure preventive and consistent-up-to-date education and guidance are accessible through an easily found self-service location. If achieved, such self-service delivery allows minimal effort for a person using the service. However, interactive updates require the attention and oversight of an individual responsible for that part of the content cycle. Agentic AI service provision can efficiently engage customers, service provision personnel, and case provision personnel, while also enabling such updating activities.




5. Orchestrated RAG Pipelines in Practice

Successful orchestration of RAG pipelines at scale requires careful consideration of multiple aspects, including source material, data ingestion and indexing, retrieval layer design, agentic decision-making policies, and autonomy boundaries. Crucial questions include which data sources will be leveraged, what retrieval layer has to be implemented, and how decision-making capabilities can be provided to the agent that operates the orchestration. Persistently and accurately answering these questions often proves to be tedious in practice and specifically requires a pragmatic design-oriented approach when applied to the governance-native agentic AI architecture for enterprise ITSM and HRSD automation considered here.

The deployment context is thus part of a larger enterprise with an IT service management (ITSM) department and a human resources service delivery (HRSD) department. In line with their respective responsibilities, ITSM serves as the operational executor of orchestration requests in the domain of business technological queries and actions, while the employee lifecycle within the enterprise is managed by HRSD. The decision policies for handling these service delivery domains therefore differ in detail, but the overall orchestration implementation remains unchanged.

5.1. Data Sources and Ingestion

In the RAG framework, content undergoes focused extraction and organization in a dedicated knowledge store to facilitate retrieval based on incoming queries. The use of multiple sources—external and corporate, structured and unstructured—enriches the RAG approach. External data sources include web pages, online documents, wikis, databases, and other platforms, while corporate data sources encompass unstructured and structured internal knowledge stores, email, wikis, service desks, and external service providers. Internal knowledge stores require attention to ensure accuracy, and nudges to perform periodic verifications throughout lifecycle phases can




be integrated into workflows.



Fig 3: Knowledge-Centric RAG Architectures: Ingest Dynamics and Governed Verification for Adaptive LLM Performance

The ingest process organizes and extracts relevant content from these data sources, enabling focused searches. A structured knowledge store ensures the availability of up-to-date information and context for accurate, responsive, and helpful interactions. Updating changes must be triggered promptly, even proactively, for significant updates to dynamic external sources such as Exchange Data and Knowledge Centers. Several ingest concepts are applicable, contingent on particular enterprise needs, capabilities, and compliance. These concepts—Managed, Auto, Nudge, and External Publish Ingest—engagement decisions are either manually defined or defined as part of governance policies, and can also vary across data source types.

5.2. Retrieval Layer Design



The retrieval layer of orchestrated RAG pipelines serves an instrumental function by determining what data will be leveraged to generate responses to end-user requests, as well as contributing indirectly to the generation process. While the pipelines' generators can be calibrated and tuned to ensure high-quality outputs, system operators remain susceptible to hallucinations, errors, and bias if the sources from which information is retrieved are poor proxies for the content being requested. Prioritizing relevant and high-fidelity sources therefore increases the likelihood of successful assistant enablement.


The nature of the sources included, their curation and preparation, and how these aspects are integrated into the retrieval layer's design for each orchestrated pipeline depend not only on the target domain or subject matter of the requests but also on the requirements of the specific enterprise function engaged by the pipeline. Such factors will arise from the requirements analysis during the orchestrated pipeline design phase. A system-oriented RAG pipeline is conducive for factual queries and when referencing predominantly structured data. A user-oriented RAG pipeline requires a set of data parameters that cover the breadth of end-user needs and address the nuances of human interaction: empathy, mitigation of frustrations, clarifications, and a naturalistic dialogue experience.

6. Agentic Decision-Making and Autonomy Boundaries


Governance-Native Agentic AI Architecture for Enterprise ITSM and HRSD Automation Using Orchestrated RAG Pipelines: Objectives, Principles, and Architectures. Agentic Capabilities and Roles. Agentic Decision-Making and Autonomy Boundaries

Automation of IT Service Management (ITSM) and Human Resource Service Delivery (HRSD) in an enterprise environment requires an operating model and technology architecture capable of supporting mission-critical services. These services necessitate an appropriate layer of supervision/oversight and risk mitigation, but cannot be efficiently delivered by a human-only organization. Enabling the efficient use of Orchestrated Retrieval-Augmented Generation (RAG) Pipelines for automation requires the establishment of a governance-native agentic AI architecture that both mitigates the risks inherent in automation and enables a high level of automation efficiency for enterprise ITSM and HRSD.

In the context of enterprises, ITSM and HRSD incident resolution by non-IT and non-HR resources is typically unpredictable, poorly documented and of low quality. During an incident, these resources require advisory and execution support, and risk/reputation mitigation measures.



These artifacts and functionality do not justify the establishment and ongoing update operation of an engineering knowledge graph. However, for known-unknown or unknown-unknown failures with large impact for which non-IT and non-HR resources require action by IT or HR resources, the operation of an automation service for advice, execution and audit trail generation is well justified. Supporting incidents with an appropriately resourced and capable automation service therefore meets the high-priority improvement objective of efficient automation of delivery, risk and reputation mitigation, and knowledge capture.




Equation 3) Converting similarity scores into retrieval probabilities (softmax)

To combine multiple retrieved chunks, map scores to probabilities:

Let $s_i = \cos(e_q, e_i)$. Then:

$$P(c_i | q) = \frac{\exp(\alpha s_i)}{\sum_{j=1}^k \exp(\alpha s_j)}$$

Step-by-step



1. Exponentiate scores so higher scores get disproportionately higher weight.
2. Divide by sum to ensure probabilities sum to 1.
3. α is a “temperature/scale” controlling sharpness.


6.1. Decision Policies and Constraints

In practice, agentic initiatives are constrained by decision policies that define the breadth of allowable decisions and actions, as well as by additional constraints that establish boundaries for language generation (e.g., specific attributes to be included or omitted from a travel itinerary). Care must be taken to ensure that decision policies do not implicitly remove substantial agentic capabilities—for example, routing travel-related requests or funds through a corporate travel office versus allowing employees to plan and fund their own travel using personal funds. Decision policies are responsible only for delineating lines of responsibility and accountability, rather than specifying what is a “right choice.” Indeed, a well-designed decision policy can be viewed as an experience-sharing mechanism, embedding lessons learned from past decisions in a way that helps to guard against known pitfalls without constraining future agents’ thinking too much.

Additional constraints ensure that pipelines are used in appropriately sensitive contexts and for appropriately sensitive purposes (e.g., a corporate travel itinerary should not be generated using a public ChatGPT API key). Natural-language generation capabilities are also supplemented where desirable and practicable, including, for example, through the use of banking-and-finance-specific content emitted by a pipeline that provides a specialized banking-and-finance-focused document corpus and accompanying structured data tables to an agent tasked with acting as a travel office.

6.2. Oversight, Interventions, and Escalation

Unfortunately, the AI systems underpinning HRSD, enterprise ITSM implementation, and other use cases retaining supervisory deference are not infallible. Incorrectly generated results may pose reputational or datacenter-security risks, unjustifiably imply certainty, or contain sensitive information sourced from the corporate data lake without consent. Oversight of all machine-generated output through appropriate patterns, suggesting human-alignment checks




for lower-confidence responses, has already been established. For increasingly low-quality machine-generated output suggesting possible degradation in competence, three remedial patterns have been identified. Roles are colour-coded for added clarity. In the first remedial pattern, busy fulfilment roles delegate fulfilment to focusing fulfilment roles but remain available for sphymic checks. In the second remedial pattern, fulfilment roles request that neutral or uninformed roles initiate perspective-taking checks. The third remedial pattern strategically escalates poorly constrained fulfilment roles to decision-management roles capable of vetoing incorrect commands and content-disclosure outputs.

As a further preventative measure, two advisory patterns check fulfilment-related topic areas. The first advisory pattern uses a sphymic perspective to validate subtle final outputs in fulfilment commands, whereas the second advisory pattern confirms that non-fact-based projects are free from holiday-related content. Finally, outputs matching trusted prompts are provided by classified fulfilment roles capable of assuming supervisory responsibility for knowledge-area integrity or security.

7. Conclusion

Recent events on a global scale have engendered popularization, democratization, and, importantly, commoditization of artificial intelligence (AI). Combined with a concurrent push by technology vendors for cloud-based service delivery models, an emerging class of ready-to-use AI products is now infiltrating the enterprise IT environment. While these systems offer the promise of transforming the business landscape, widespread adoption also presents serious risks and challenges. In particular, problems associated with bias, accuracy, hallucination, intellectual property infringement, and privacy violations are well-documented. Vargas et al. argue that these issues must be addressed before organizations deploy such systems in mission- or enterprise-critical functions. Holistically governing the AI solutions being exploited or developed at various levels within the business is pivotal to meeting these requirements. The Gan.ai2 framework articulates governance in terms of models and operating principles, as well as embedding governance into the tasking, deployment, and usage of AI systems.

Following this approach, a governance-native architecture designed expressly for agentic implementation is presented. Informed by the activity domains of enterprise information technology service management (ITSM) and human resource service delivery (HRSD), the architecture is capable of realizing service fulfilment needs through orchestrated retrieval-augmented generation (RAG) pipelines. The particular characteristics of agentic systems afford




additional levels of service delivery speed, consistency, and volume. Demonstration of the architecture in practice elucidates its ability to deliver on these expectations. These aspects therefore jointly highlight an enterprise-context model of governance-native agentic AI defined in Vargas et al. and the significance of the supported service environment to the broader ecosystem.



Fig 4: Agentic Service Delivery Performance

7.1. Final Reflections and Future Directions

In recent years, Generative Large Language Models (GLLMs) have enabled enterprise agility and business productivity by making subject matter knowledge readily and conveniently accessible in natural language. The emergence of Artificial Intelligence (AI) technologies with the capability of functioning as agents that directly interact with APIs, database applications, and




third-party services is poised to revolutionize digital operations. In such scenarios, AI has the potential to complete tasks with minimal user intervention or even entirely autonomously. However, task-oriented AI in critical domains with risk of considerable harm to human life or the environment requires governance-active, intelligence-based safeguards to mitigate risk. During the 1970s and 1980s, Herbert Simon elucidated the nature of Sovereign Intelligence and proposed the establishment of a Societal Intelligence Framework with active oversight for IT applications able to learn, adapt, and operate without direct user intervention.

The AI governance frameworks currently mandated by national or regional regulators are fundamentally addressing bias risk and fail-safety. These aspects represent just two of the many features distinguishing agents from more basic examples of automating systems. A Framework for Governance-Native Agentic AI architecture that embraces key concepts of Governance typically articulated in corporate, national or regional Governance domains has been developed and applied in the context of Enterprise Information Technology Service Management (ITSM) and Human Resource Service Delivery (HRSD) within a Global 500 multinational corporation. The Architecture is organized around Generative Retrieval-Augmented Generation (RAG) Pipelines, where data and AI support for Processes and Tasks are designed for distinct Structured, Adaptive or Instinctual Processes and supportive Task Types.

References

1. Al-Alawi, A. I., Naureen, M., AlAlawi, E. I., & Naser Al-Hadad, A. A. (2021). The role of artificial intelligence in recruitment process decision-making. In 2021 International Conference on Decision Aid Sciences and Application (DASA) (pp. 197–203).
2. Almeida, P. G. R. de, dos Santos, C. D., & Farias, J. S. (2021). Artificial intelligence regulation: A framework for governance. *Ethics and Information Technology*, 23, 505–519.
3. Segireddy, A. R. (2020). Cloud Migration Strategies for High-Volume Financial Messaging Systems.
4. Budhwar, P., Malik, A., De Silva, M. T., & Thevisuthan, P. (2022). Artificial intelligence—challenges and opportunities for international HRM: A review and research agenda. *The International Journal of Human Resource Management*, 33(6), 1065–1097.
5. Chhillar, D., & Aguilera, R. V. (2022). An eye for artificial intelligence: Insights into the governance of artificial intelligence and vision for future research. *Business & Society*, 61(5), 1199–1231.
6. Erdélyi, O. J., & Goldsmith, J. (2022). Regulating artificial intelligence: Proposal for a global solution. *Government Information Quarterly*, 39(4), 101748.



7. Gahnberg, C. (2021). What rules? Framing the governance of artificial agency. *Policy and Society*, 40(2), 194–210.
8. Nagabhyru, K. C. (2022). Bridging Traditional ETL Pipelines with AI Enhanced Data Workflows: Foundations of Intelligent Automation in Data Engineering. Available at SSRN 5505199.
9. Gill, A. S., & Germann, S. (2022). Conceptual and normative approaches to AI governance for a global digital ecosystem supportive of the UN Sustainable Development Goals (SDGs). *AI and Ethics*, 2, 293–301.
10. Gong, Y., Zhao, M., Wang, Q., & Lv, Z. (2022). Design and interactive performance of human resource management system based on artificial intelligence. *PLOS ONE*, 17(1), e0262398.
11. 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.
12. Mäntymäki, M., Minkinen, M., Birkstedt, T., & Viljanen, M. (2022). Putting AI governance on the agenda: A systematic literature review of AI governance. *Internet Research*, 32(7), 1–31.
13. Palos-Sánchez, P. R., Baena-Luna, P., Badicu, A., & Infante-Moro, J. C. (2022). Artificial intelligence and human resources management: A bibliometric analysis. *Applied Artificial Intelligence*, 36(1), 2145631.
14. Papagiannidis, E., Enholm, I. M., Dremel, C., Mikalef, P., & Krogstie, J. (2022). Toward AI governance: Identifying best practices and potential barriers and outcomes. *Information Systems Frontiers*, 25, 123–141.
15. Gottimukkala, V. R. R. (2020). Energy-Efficient Design Patterns for Large-Scale Banking Applications Deployed on AWS Cloud. *International Journal of Advanced Research in Computer and Communication Engineering (IJARCCE)*, DOI, 10.
16. Pan, Y., Froese, F., Liu, N., Hu, Y., & Ye, M. (2022). The adoption of artificial intelligence in employee recruitment: The influence of contextual factors. *The International Journal of Human Resource Management*, 33(6), 1125–1147.
17. Qamar, Y., Agrawal, R. K., Samad, T. A., & Jabbour, C. J. C. (2021). When technology meets people: The interplay of artificial intelligence and human resource management. *Journal of Enterprise Information Management*, 34(5), 1339–1370.
18. Radu, R. (2021). Steering the governance of artificial intelligence: National strategies in perspective. *Policy and Society*, 40(2), 178–193.



19. Rastogi, S., & Sharma, P. (2022). Revisiting artificial intelligence in HRM. *SJCC Management Research Review*, 12(2), 110–128.
20. Aitha, A. R. (2021). Dev Ops Driven Digital Transformation: Accelerating Innovation In The Insurance Industry. *Journal of International Crisis and Risk Communication Research*.
21. Schneider, J., Abraham, R., Meske, C., & vom Brocke, J. (2022). Artificial intelligence governance for businesses. *Information Systems Management*, 40(3), 229–249.
22. Sithambaram, R. A., & Tajudeen, F. P. (2022). Impact of artificial intelligence in human resource management: A qualitative study in the Malaysian context. *Asia Pacific Journal of Human Resources*. Advance online publication.
23. Stix, C. (2022). Foundations for the future: Institution building for the purpose of artificial intelligence governance. *AI and Ethics*, 2, 463–476.
24. Taeihagh, A. (2021). Governance of artificial intelligence. *Policy and Society*, 40(2), 137–157.
25. Ulnicane, I. (2021). Framing governance for a contested emerging technology: Insights from AI policy. *Policy and Society*, 40(2), 158–177.
26. Vrontis, D., Christofi, M., Pereira, V., Tarba, S., Makrides, A., & Trichina, E. (2022). Artificial intelligence, robotics, advanced technologies and human resource management: A systematic review. *The International Journal of Human Resource Management*, 33(6), 1237–1266.
27. Inala, R. Designing Scalable Technology Architectures for Customer Data in Group Insurance and Investment Platforms.
28. Sreerangapuri, A. (2021). IT service management transformation: Implementing AIOps for proactive service delivery. *NeuroQuantology*, 19(10), 147–155.
29. Yang, C., Chen, Y., & Li, X. (2021). Research framework for determining how artificial intelligence enables information technology service management for business model resilience. *Sustainability*, 13(20), 11496.
30. Abonamah, A., Torre, D., Poulin, M., & Repetto, M. (2022). Explainable artificial intelligence in human resources: A computational study. In *2022 International Conference on Data Analytics for Business and Industry (ICDABI)* (pp. 1–6).
31. Choudhury, P., Foroughi, C., & Larson, B. (2021). Work-from-anywhere: The productivity effects of geographic flexibility. *Strategic Management Journal*, 42(4), 655–683.