ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

Orchestrating LLMs and SLMs for Autonomous Enterprise Service Automation

Ethan Williams
Asst Professor

Abstract

Autonomous enterprise agents designed for orchestrated decision automation across information technology service management (ITSM), human resources service delivery (HRSD), and customer service and support (CSS) platforms are detailed. Growing enterprise demand for automation and decision support hinges on self-governing software agents capable of independent task scheduling and execution, yet ITSM, HRSD, and CSS environments present complex governance requirements that hinder the deployment of agent-based systems. Empirical evidence points to the existence of a shared set of conditions—standardized data models and commonplace processes, together with solutions supporting rudimentary automation—beneath the apparent diversity of cross-domain decision workflows. An agent composition is described, with an orchestration layer addressing scalability and risk management.

The software agents underpinning enterprise automation seek to relieve staff of routine, low-level decisions while delivering consistent, reliable results. Orchestrated instances feature additional layers facilitating expansive collaboration through the decomposition of complex tasks into multiple subtasks and the admission of new capabilities via third-party extensions. Two intertwining sets of requirements reflect the diversity of enterprise domains: the specific conditions imposed by service management platforms—issues governing service-catalog operations, ticket lifecycles, and incident-resolution processes—and wider concerns of privacy, confidentiality, and data protection, which are paramount when handling employee information.

Keywords: Autonomous Enterprise Agents, Decision Automation Systems, IT Service Management Automation, Human Resources Service Delivery, Customer Service And Support Platforms, Agent-Based Orchestration, Self-Governing Software Agents, Cross-Domain Decision Workflows, Standardized Data Models, Enterprise Process Automation, Orchestration Layer Design, Scalability And Risk Management, Task Decomposition Frameworks, Collaborative Agent Systems, Third-Party Capability Extensions, Service Catalog Operations, Ticket Lifecycle Management, Incident Resolution Automation, Privacy And Confidentiality Requirements, Enterprise Data Protection.

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

1. Introduction

Enterprises seek to automate repetitive tasks to free people for higher-value activities. The growing adoption of Service Management tooling enables automation of a subset of these tasks by integrating systems across the enterprise. IT Service Management (ITSM) systems provide a mechanism to define standardized service catalogs, lifecycle management workflows and governance controls. Most common service requests involve provisioning of new IT resources (laptops, accounts, cloud resources) or obtaining support for broken ones. They are relatively similar for all organizations operating in the same industry. Therefore, decision automation of these requests can be achieved at scale by sharing a common data model, policy definition, decision-making body, skill requirements and audit data for all organizations. Service categorization is the primary decision that defines the request patterns for an organization and their planning and execution.

When these repetitive processes span multiple functional areas like IT, HR and Customer Support, their automation is not yet genuinely enterprise scale. Business Operation Tasks corresponding to the Employee Lifecycle and Customer Interaction Lifecycle Management processes exhibit high volumes of cross-domain HRSD and CSM handoffs. These connections are often non-standard and therefore difficult to automate. Recent advances in Large Language Models open up possibilities to redefine these processes as workflows that are rapidly translatable into automation scripts for various tools like ServiceNow. Holistic consistency of Service Operations can also be achieved by orchestrating autonomous enterprise agents that are responsible for automating repetitive business operations. These autonomous enterprise agents can operate independently and across organizational boundaries independently on behalf of every enterprise.

1.1. Overview and Significance

Automation of Enterprise Processes with Distributed Decision Automation by Orchestrated Autonomous Agents explores orchestration of agents that assist human decision making in a Service Management Environment—complementing the earlier focus on enterprise orchestration of decision automation processes, and breaking new ground on autonomous agents that presume Enterprise models and governance. Enterprises collaborate in Service Management Environments represented by three domains—Information Technology (ITSM: IT Service Management), Human Resources (HRSD: Human Resource Service Delivery), and Customer

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

Service (CSM: Customer Service Management). These domains encompass operational processes centralized on service delivery, multistep task execution, and demand-response ticket communications. Service delivery processes can lead to hybrid scenarios based on repeated, common decision patterns spanning Process Areas. Distribution of these decision-making tasks for automated execution by Orchestrated Autonomous Agents relieves human experts from burden without removing them from control, enabling enterprise automation at speed and scale. Quantified parameterization of Service Management Environment demand feeds the automation to deliver information at speed for consumer Contact Centres using diverse messaging channels.

Unlike Enterprise Orchestration Automation, where the Enterprise is utilizing its capabilities to automate processes and deliver services, Operational decision-making is assumed to be conducted mainly by humans because of complexity and the dynamics of the task inputs. Orchestrated Autonomous Agents support human decision-making by applying learned, common, repeatable decision patterns. Such Orchestrated Autonomous Agents automate individual tasks and manage the agent ecosystem enabling scheduling of these tasks for execution. The knowledge and data base for the Orchestrated Autonomous Agents consists of the Data Pools and Data Stores defined in the Service Management Environment representation, whether Service Management Environment-based or derived from other sectors. Successful Orchestrated Autonomous Agent execution remains limited to process areas and decision patterns deemed suitable for autonomous execution by human experts.

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

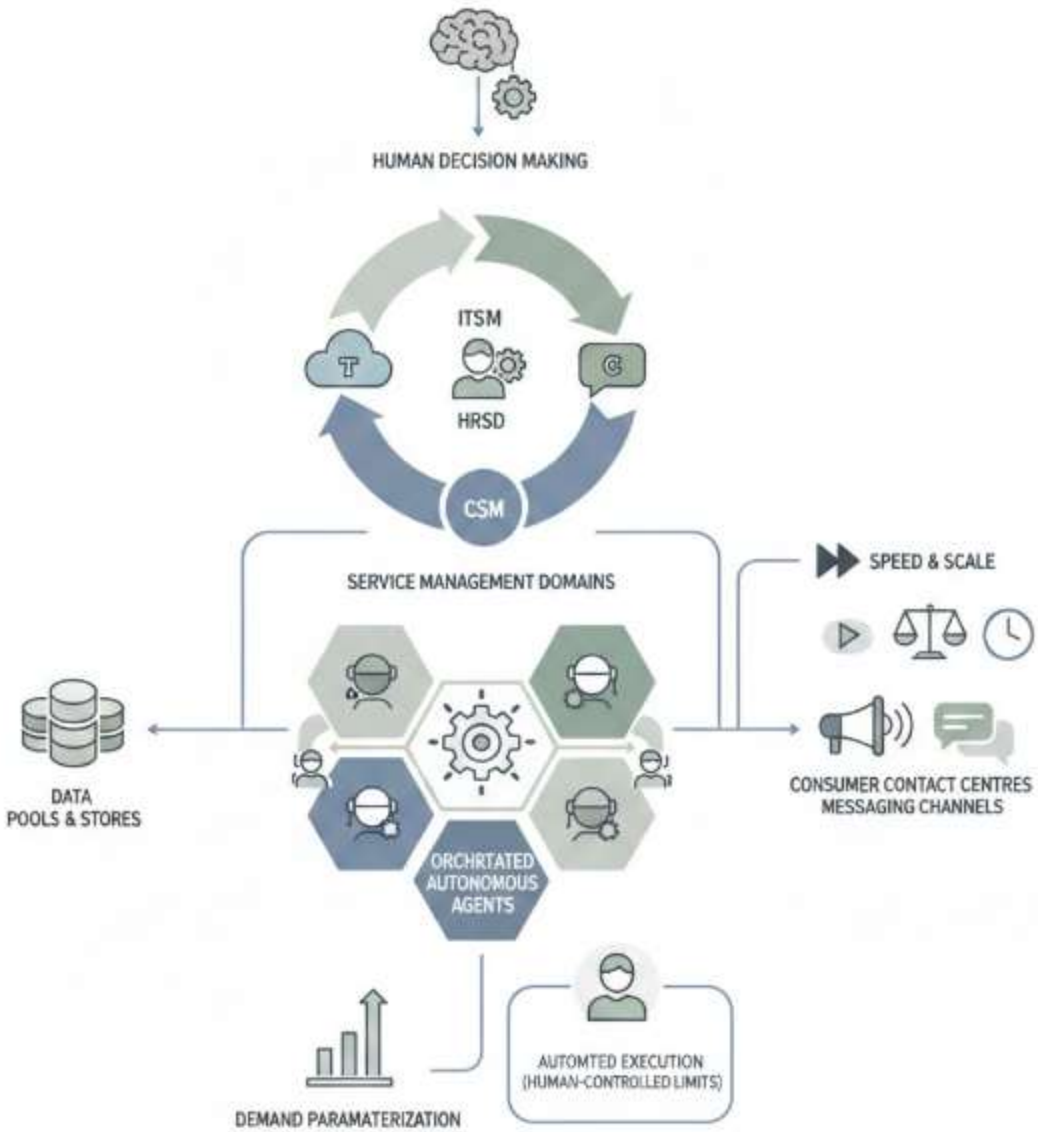

Fig 1: Orchestrating Autonomy: A Framework for Distributed Decision Automation within Multi-Domain Service Management Environments

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

2. Background and Context

For external orchestrated agents to be effective, enterprise processes must be partitioned into well-understood components. Enterprise policies, such as those found in Configuration Management Databases (CMDBs), Knowledge Bases (KBs), and Security Policy Databases (SPDs), must also be expressed in a form that supports machine comprehension. Enterprise IT Service Management (ITSM) platforms and Human Resource Service Delivery (HRSD) platforms can support employee services in a variety of ways, with orchestration facilitating the management of end-to-end service delivery and providing employees with a single view of the services to which they are entitled at any point in time. In addition to ITSM and HRSD, Customer Service Management (CSM) also offers high potential for orchestration and collaboration, especially with ITSM for Incident Management tickets and HRSD for areas such as Employee Lifecycle Services.

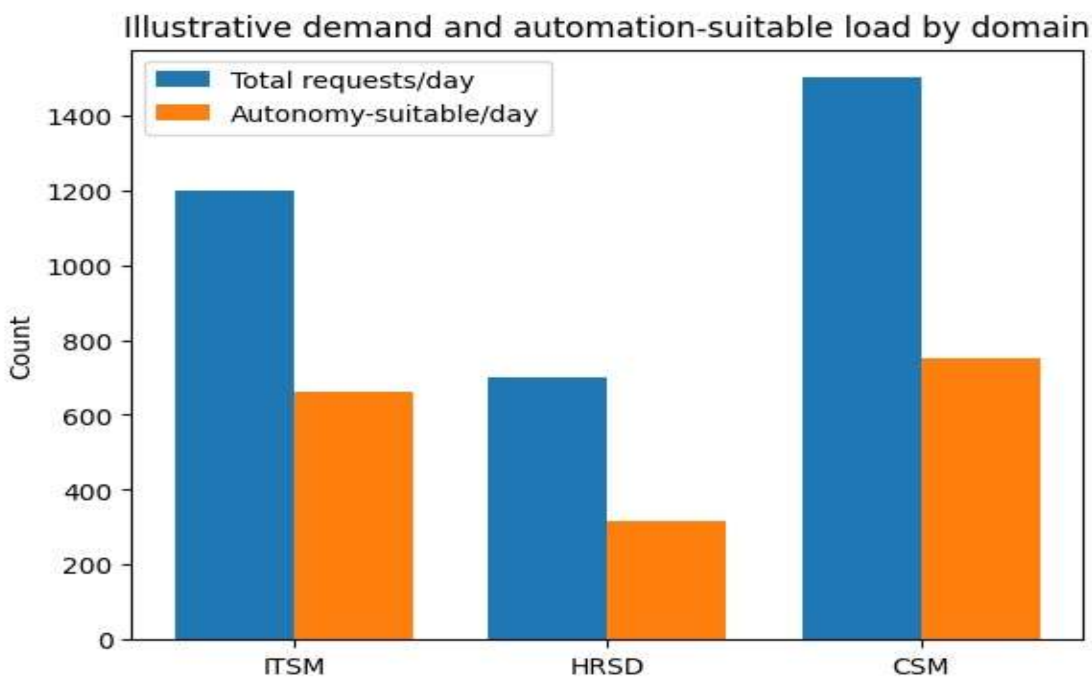

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

While information systems are often designed for enterprise automation, enterprises often rely on humans to make decisions during service execution. Orchestrated agents can help transform enterprise information systems from automation-supported systems into enterprise automation systems. Orchestrated agents can also support enterprise agent-based decision automation by empowering autonomous agents, in contrast to simply providing automation for autonomous agents. Processes under decision boundaries can be orchestrated into a top-down process and end-to-end decision and automation processes can be defined to accommodate common criteria across service delivery domains.

Equation 1: Step-by-step equations for task decomposition (DAG model)

1.1 Build a dependency graph

Let the task be T. Decompose into n subtasks:

1. **Subtask set**

$$\mathcal{V} = \{1,2,\dots,n\}$$

2. **Directed edges (dependencies)**
   If subtask i must finish before j starts, add edge $(i \to j)$.

$$\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$$

3. **DAG**

$$G = (\mathcal{V}, \mathcal{E})$$

1.2 Valid schedule constraints

Let s_i be start time of task i, and d_i its duration.

For each dependency $(i \to j) \in \mathcal{E}$:

$$s_j \geq s_i + d_i$$

2.1. Foundations of Autonomous Enterprise Agents

ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

The concept of autonomous enterprise agents is rooted in the theory of multi-agent systems and is distinguished from well-known idées de recherche such as process orchestration, Robotic Process Automation (RPA), and Artificial Intelligence. A specific type of agent combines characteristic features of all of these trend-setting ideas and is shown to form an effective technology for practical enterprise applications, where RPA alternatively reaches its limits. The research is substantiated by foundational publications that characterize such autonomous enterprise agents and provide the basis for addressable and governance-relevant operational requirements. The concept is further substantiated in specific application domains requiring distinct technologies for Service Management, Human Resource Service Delivery, and Customer Service Management. Both the discovered requirements and the operational concept based on autonomous enterprise agents serve as foundations for proof-of-concept implementations.

The constructed autonomous-agent-based platform design automates individual decisions and their broader supporting processes in a productive multi-agent system. Three operational layers are instantiated: The functional decomposition of decision automation, the interaction within dedicated agent organizations, and the orchestration of operational processes. The bottom layer lays out practical task and workflow patterns for automating enterprise decisions. The middle layer specifies the essential interaction patterns for specific organizational roles supporting infrastructure decisions and holistically governing the rules used for such decisions. The top orchestration layer describes the overall coordination of decision automation within larger enterprise environments, fulfilling external policy requirements, and supporting both scalability and reliability aspects.

Equation 2: Step-by-step equations for agent composition under budget–quality trade-offs

2.1 Candidate agent pool

Let $\mathcal{A}$ be the set of available agents (some may use LLMs, some SLMs).

For each subtask i and agent a:

- $c_{i,a}$ = cost to have agent a do subtask i (money/time/compute)
- $q_{i,a} \in [0,1]$ = expected quality/success probability

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

- $t_{i,a}$ = expected duration

Decision variable:

$$x_{i,a} = \begin{cases} 1 & \text{if agent } a \text{ assigned to task } i \\ 0 & \text{otherwise} \end{cases}$$

2.2 Assignment constraints (each task gets one agent)

$$\sum_{a \in \mathcal{A}} x_{i,a} = 1 \quad \forall i \in \mathcal{V}$$

2.3 Budget constraint

Let total budget be B.

$$\sum_{i \in \mathcal{V}} \sum_{a \in \mathcal{A}} c_{i,a}\, x_{i,a} \le B$$

2.4 Quality objective (maximize expected plan quality)

A simple (and common) objective is maximize sum of expected qualities:

$$\max \sum_{i \in \mathcal{V}} \sum_{a \in \mathcal{A}} q_{i,a}\, x_{i,a}$$

This directly implements “modeling budget and quality considerations” .

2.5 Optional: multi-objective (budget vs quality)

Sometimes you want a single scalar objective:

$$\max \sum_{i,a} q_{i,a}\, x_{i,a} - \lambda \sum_{i,a} c_{i,a}\, x_{i,a}$$

Where $\lambda \ge 0$ is a trade-off weight (higher $\lambda \to$ cheaper, lower quality mix).

ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

2.2. ITSM, HRSD, and CSM: Requirements for Automation

The decision-making processes and supporting data models commonly found in IT Service Management (ITSM), Human Resource Service Delivery (HRSD), and Customer Service Management (CSM) environments possess sufficient structure and organization to facilitate automation that is not simply focused on executing predefined responses. Although the identified workflows frequently lack the volume, reliability, or predictability that training-based AI models require to approach human-like performance, they often comprise discretized decisions that can be tackled more efficiently than with human effort.

To provide transparency and auditability, the governance of these processes tends to be more formalized than in domains where automation relies on AI techniques. Such governance also supports the definition of policies that can designate the decision boundaries for the deployment of autonomous agents. Processes common to the ITSM, HRSD, and CSM domains can therefore be broken down into a series of sub-decisions, which are stitched together into a single execution plan that encompasses the complete task. Despite the potential complexity of these decision flows, the significance of the sub-decisions and their relative weights can be easily assessed and compared across agent instances. In this context, the orchestration mechanism can stabilize these decision processes against occasional mistakes while providing guaranteed correctness and fulfillability of the overall work.

Domain	Requests/day (illustrative)	Autonomy-suitable %	Human-in-loop %
ITSM	1200	55.00000000000001	20.0
HRSD	700	45.0	30.0
CSM	1500	50.0	25.0

Equation 3: Step-by-step equations for policy enforcement and decision bounds

3.1 Policy constraints as predicates

Let $\pi_k(\cdot)$ be policy k (privacy, approval thresholds, data-access limits).

Define a feasibility indicator:

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

$$\text{Feasible}(r,\text{plan}) = \prod_{k} \mathbf{1}\{\pi_k(r,\text{plan}) = \text{true}\}$$

3.2 Upper/lower bounds enforcement

Example: "risk score" of plan must lie in bounds:

- Let $R(\text{plan}) \in [0,1]$
- bounds $R_{\min} \leq R(\text{plan}) \leq R_{\max}$

$$R_{\min} \leq R(\text{plan}) \leq R_{\max}$$

This encodes "preventing upper or lower decision bounds from being violated" .

3.3 Human-in-the-loop gate as a constraint

If policy requires approval when risk is high:

$$\text{If } R(\text{plan}) \geq \tau, \quad \text{then require approval event } H = 1$$

3. Architecture of Orchestrated Agent Systems

An orchestrated autonomous enterprise agent system comprises multiple agents cooperating to achieve a shared objective, within predefined policies, while processing service requests or service-related data originating from enterprise platforms or their users. Every service category identified in the ITSM, HRSD, or CSM domain mandates agent systems providing different agent roles and capabilities. Each agent type specializes in discrete responsibilities—request initiation, planning, task execution, or monitoring—decomposing intricate activities into manageable elements.

Workflow diagrams illustrate scenarios involving three dedicated service catalogs; connections with primary platforms are shown for clarity. Artefacts spanning the agent software deployment, interfaces with platforms or users, and associated data flows are listed. The orchestration layer coordinates the services, synchronizing activity across smaller systems; it

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

implements policies transversing multiple domains and domains with higher criticality, specialty reserves, or elevated approval criteria. Security aspects safeguard sensitive information; mechanisms detect and recover from internal errors, service interruptions, and other operational disruptions.

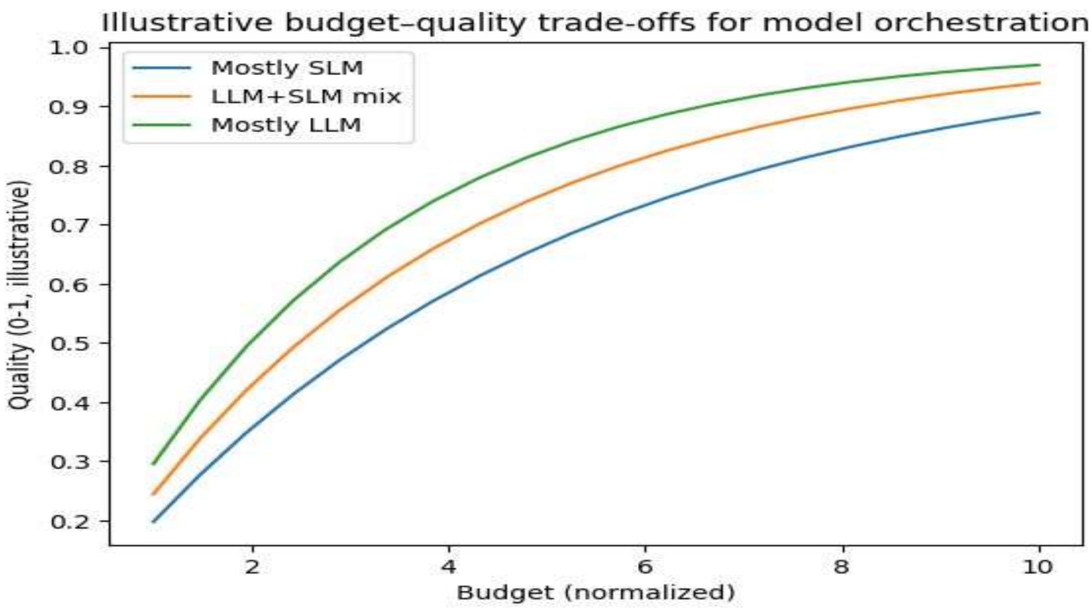

Agent role	Primary output artifact
Intake/Router agent	Service category + context
Planner agent	Task DAG + schedule
Executor agent	Task results + ticket updates
Policy/Governance agent	Approve/deny + constraints
Monitor/Recovery agent	Alerts + rollback triggers

3.1. Agent Composition and Roles

ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

Enterprise automation provides an opportunity to move away from the traditional model of direct human involvement in operational tasks. Rather than delegating tasks to IT personnel, employees can fulfill their responsibilities by directly requesting these services through a service catalog that specifies the processes governing the tasks to be fulfilled. Some of those service requests can be fulfilled automatically by automated decision actors that are endowed with a partial understanding of the enterprise and its context. Their understanding is partial in that they need not have complete knowledge of the enterprise in order to execute decisions, relying instead on information for particular decisions that comes from other services.

The approach described here involves multiple-level orchestration among decision actors and the ITSM, HRSD, and CSM platforms to enable scalable automation of tasks and decisions. Processes execute inline when needed by a platform (e.g., planning an onboarding process as part of the incident management workflow). Service requests from the enterprise for cross-domain services, such as employee lifecycle processes or leave-of-absence requests, invoke the planning and execution of service workflows.

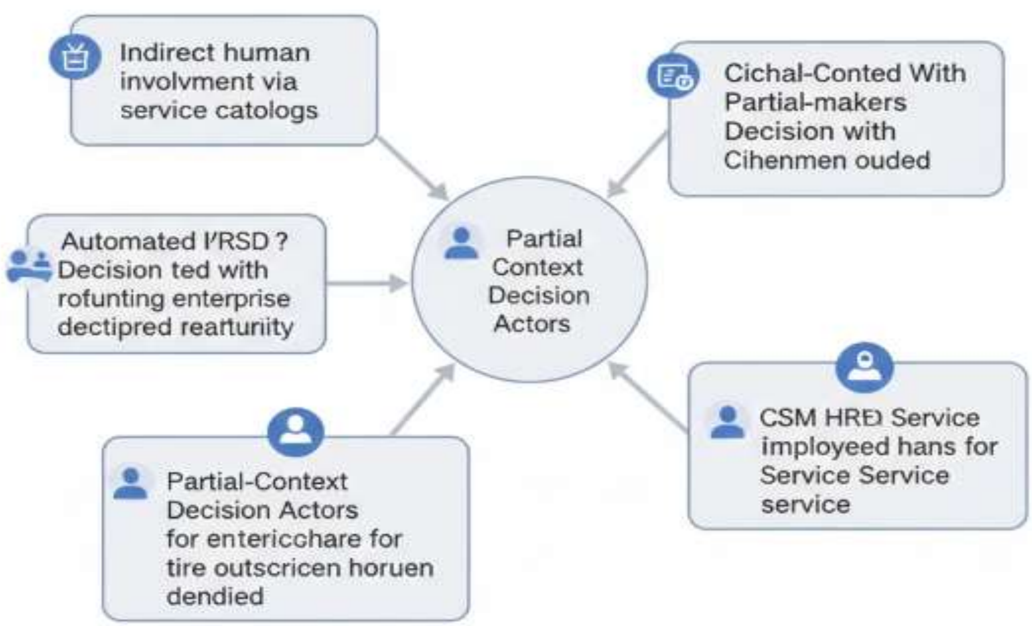

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

Fig 2: Multi-Level Orchestration of Partial-Knowledge Decision Actors: A Cross-Domain Framework for Scalable Enterprise Automation

3.2. Orchestration Layer: Coordination and Policy Enforcement

The orchestration layer coordinates the activities of composed agent systems and enforces service-level governance policies without involving the core decision-making functions of task execution agents. Policies are implemented through a dedicated agent responsible for analysis and cross-agent reasoning that can exploit higher-level strategic objectives or operational constraints, enhancing enterprise decision provenance and risk mitigation by preventing upper or lower decision bounds from being violated.

Service automation candidates are either actively or passively involved in incident management systems, relying on an information model to automatically assess candidate service maps and execution capacity under specific policies. Clustered incident handling by service automation agents allows scale by managing multiple related requests as a single task, with traditional ITSM roles acting as resource providers. Direct CSM integration facilitates human collaboration within services or processes of connected consumers, including situations with shared infractions across consumers serviced by different organizations. Employee-lifecycle-related services are either provided by HR systems or orchestrated when external and internal sessions need to be linked to facilitate support interactions during requests conveyed through CSM channels. Access-control characteristics shape the granularity of integrated operation modules, data mapping, and offered services depending on employees' lifecycle and consent.

4. Decision Automation Processes

Orchestrated enterprise systems help automate enterprise decisions based on processes that can easily be understood and adequately expressed in business terms. Decision automation is distinct from classical process automation via information technology and represents the expansion of an automated enterprise's automated operational processes into areas where the end result can still be expressed and therefore formalized with a reasonable degree of certainty. Decision automation enables an enterprise to make timely, transparent, and repeatable decisions based on the conditions revealed by integrated data, in contrast to traditional human decision-making processes that tend to operate in silos and thereby can introduce biases detrimental to the enterprise's success.

ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

The key components of decision automation are the decomposition of a task or decision into a series of sub-tasks or dependencies, the planning and scheduling of the decision using a standard library of decisions, the validation against express business constraints, the execution of the decision, and the accessibility to a proven and auditable operational provenance. For any enterprise, including IT service management (ITSM), human resource service delivery (HRSD), and customer service management (CSM), the decision automation processes allow decisions to be made at the appropriate boundaries, thereby enabling improved efficiency and scalability.

4.1. Task Decomposition and Planning

Effective automation of complex tasks typically involves decomposing them into sub-tasks, assigning responsibility to (possibly different) agents, and producing a detailed plan of actions together with their execution schedule. For enterprise-scale task automation, the set of composing agents is selected on a per-run basis, dynamically determined by modeling budget and quality considerations. Task composition, together with the production of a detailed plan and schedule, is then the focus of the orchestration layer.

The rules for a standardized orchestration layer specify which responsibilities need to be covered when composing agents for the execution of particular tasks, as well as possible budget and quality trade-offs. An orchestrated agent system then implements these rules and combines output from different agents in the background to generate a cohesive view of the mission, which is used to setup the necessary decision automation flows. The definition of composing agents is aligned with the standard service catalog from which it derives all necessary inputs, ensuring a seamless operational integration between direct human interactions and those executed in the background. This integration results in a smooth and unobtrusive operation of the autonomous agents coordinating the enterprise decision automation process.

5. Interaction with ITSM, HRSD, and CSM Platforms

Increased decision automation requires integration and interaction with ITSM, HRSD, and CSM platforms. Consider touchpoints in service catalogs, incident management workflows, ticket lifecycles, employee lifecycle services, HR workflows, and cross-domain handoffs.

Service catalogs expose automation capabilities to employees and facilitate touchpoints to incrementally resolve cross-domain incidents without human intervention. The catalogs map to enterprise service catalogs and include automated incident management workflows that create

ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

system-level tickets on behalf of employees. Cross-domain inquiries and transactions trigger automated human-in-the-loop decision processes that satisfy authorization, privacy, and compliance requirements. Result modifications require transparent approval during human-in-the-loop processing.

Integrations with ITSM, HRSD, and CSM platforms allow orchestration of employee lifecycle services like onboarding, offboarding, and promotion while spanning common touchpoints like parental and medical leaves. Automated handling of common employee lifecycle services shifts resources toward complex, uncommon requests. Authorized cross-platform data access informs enterprise HR domain workflows for PTO, promotion, parental, and other enterprise services or agreements. Access controls on private data and team memberships further mitigate risk and align processing with privacy governance.

5.1. Service Catalogs and Incident Management

Service catalogs, supported by catalogs of standard services, applications, and components, provide a framework for managing service requests and incidents. By mapping these processes to an orchestrated agent system, automation opportunities can be identified. Service catalogs describe the services offered by an organization to the employees, customers, or partners. Employees use the catalog of the internal IT department to raise requests for standard services such as a new laptop or the provisioning of an application. External customers use the catalog of the customer service operations to raise service requests such as a request for credit card replacement or the raising of an incident. The customer service platform maps a request to an applicable service line or incident management workflow. An orchestrated agent autonomously controls standard requests from service catalogs that are well understood with deterministic compute requirements and clear objectives.

Some employee lifecycle services cross Human Resource Service Delivery and IT Service Management domains, requiring cooperation across teams to fulfill requests. Such cross-team dependencies are commonly managed by HRSD platforms. An orchestrated agent identifies the operation from the service catalog, orchestrates the decision-making process across various domains, triggers the required workflows in HRSD and ITSM platforms, and monitors progress until the task is completed. Employees are kept updated on the service request status using the appropriate communication channels. The orchestrated agent manages such services according to internal privacy policies, helps in compliance requirements of industries such as financial

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

services, and ensures the service request is completed only when all relevant nodes are provided. Privacy policies ensure that elements such as source account details remain hidden and are not provided in the same operation to avoid any malicious usage.

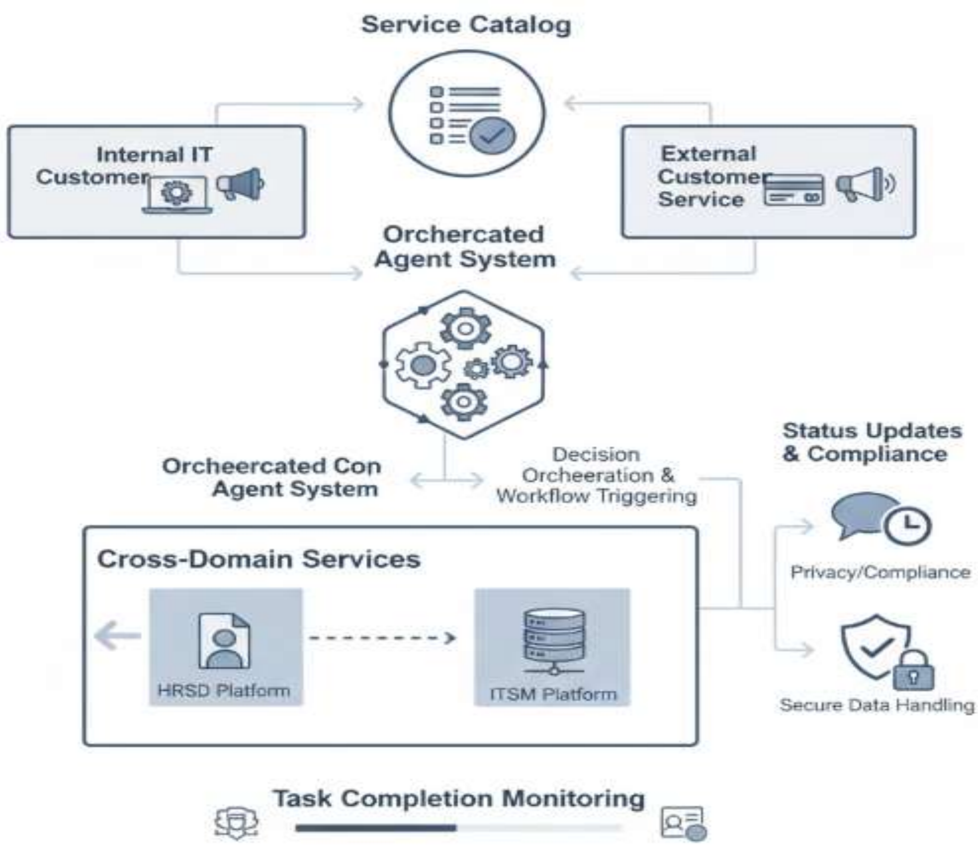

Fig 3: Autonomous Governance of Cross-Domain Service Catalogues: Orchestrating Agent-Driven Compliance and Life-Cycle Fulfillment

5.2. Employee Lifecycle Services and HR Workflows

Autonomous agents in enterprise settings must support end-to-end management of scenarios involving human employees in their capacity as employees. HR subdomains must

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

therefore have supply and demand sides, with a governance framework ensuring that agents cannot act in violation of privacy, compliance, and other policies. These requirements lead to the construction of an HR service catalogue that is governed by an enterprise privacy policy defining which data about an employee whose employee lifecycle is being managed a system actor is allowed to access.

An operationalisation of the Employee lifecycle service and the ENTERPRISE-AGENT orchestrating the management of Employee lifecycle processes in an ITSM context confirms that a full range of Decision Automated Tasks extending across Demand Planning, Supply Planning, Demand Fulfillment Task Management and Execution, and actually typical for employee lifecycle processes can be modelled. These processes include Employee Onboarding, Employee Position Change, Employee Offboarding, Contractor Engagement and Offboarding. Furthermore, the automation of ITSM processes to date, while including processes involving external customers, has focused on the resolution of tickets initiated by internal customers requesting IT services.

6. Challenges and Risk Mitigation

Decision automation using orchestrated enterprise agents introduces risks that must be carefully managed to ensure reliable, trusted, and secure operation. The success of intelligent assistants hinges on availability, performance, transparency, and explainability. Automated actions must not only be consistent with established governance policies but should also be understandable by the stakeholders impacted by them. If the decision-making processes are perceived as unpredictable, the trust placed in the end-to-end automation of cross-domain processes may suffer.

A multi-step validation procedure helps mitigate the risks arising from automated decisions. To achieve a reliable and trustworthy composition, automated decisions should be subjected to reasoning validation, process simulation, user feedback collection, and process execution in non-production environments. Meaningful metrics need to be established to assess the reliability of orchestration, and detection mechanisms should be developed for particular categories of failure, such as security policy violations, possible biases in machine learning models, or lack of completeness in the domain knowledge. Documentation of validation results and an adequate governance model are also required. Finally, a framework for real-time

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

monitoring is necessary, covering all aspects of risk that cannot be completely addressed through the validation pipeline.

6.1. Reliability, Transparency, and Explainability

Service catalogs in IT service management serve as central repositories for service definitions, configurations, and interfaces. Automation nodes connect and integrate into service catalogs via custom or standard connectors. Each connected ticketing service is modeled with the relevant catalog information, providing the contextual references that automate decision-evaluation and task-execution processes. Ticket lifecycles within any incident-management workflow are enhanced with specialized automation steps that read the ticket context, generate task sets, and invoke ticket-status changes based on the inferred completion status.

Where employee-lifecycle services exist, privacy and compliance settings determine which HR and IT products are applicable. When new employee lifecycles are created, flows connect to the HRSD product team as the source. When any onboarding Flow is marked completed, it triggers a cross-domain handover, setting HRSD to be the source and ITSM as the target. All handover data are securely passed, with any required data transformations, auto-generated request tickets in the target domain, and status-monitoring tasks created.

7. Conclusion

Efforts to evolve enterprise applications towards independence from user interaction and decision Calvinization processes have produced considerable results in making the operation of such applications and systems like ITSM, HRSD, and CSM configurable. It has been demonstrated that systems can manage highly repetitive incident record generation or service request creation processes independently and autonomously. These efforts covered a significant area of the enterprise towards becoming an Autonomous Enterprise. Specifically, it has been shown that ITSM platforms detect and create automation for a large amount of incident management workflows. More frequent services carried out by the HRSD product lines are largely reduced for their HR teams through the management of employee life-cycles and employee-life-cycle-related requests. Standard service requests defined by HRSD platforms can now be potentially fulfilled by those systems directly, with more complex workflows and new types of requests passed across to HR teams.

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

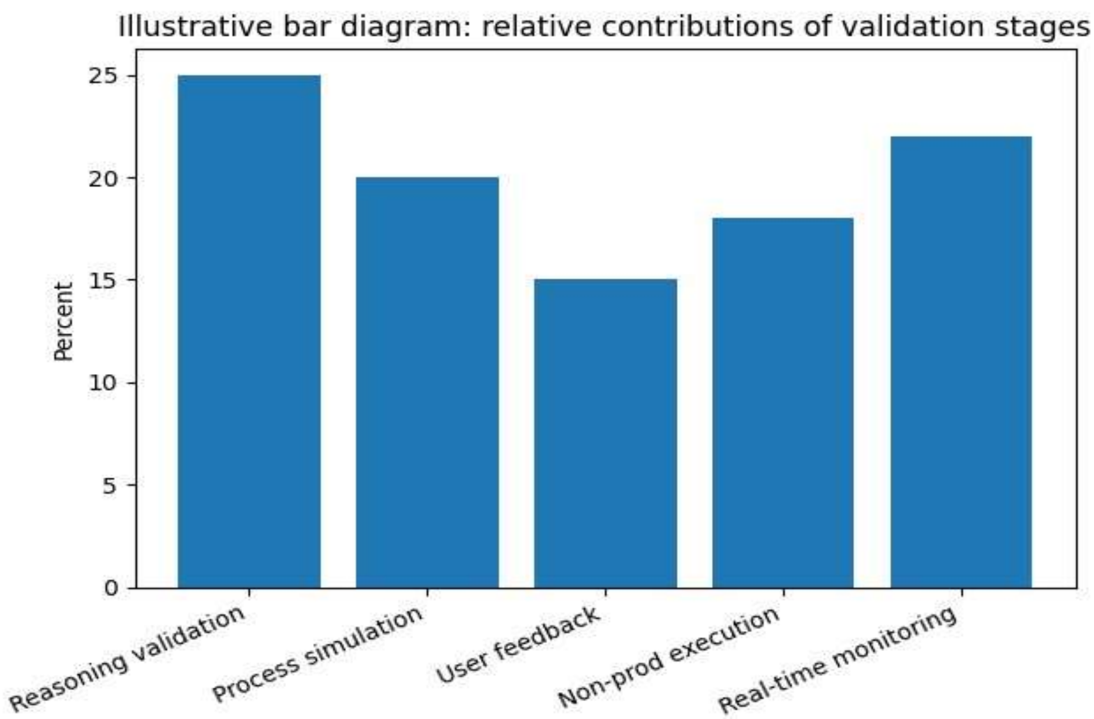

Equation 4: Step-by-step equations for reliability + the validation pipeline

4.1 Stage-wise failure reduction model

Let baseline failure probability be p_0.

Let each stage k reduce remaining failure probability by factor $(1-\alpha_k)$, where $\alpha_k \in [0,1]$ is effectiveness.

After stage 1:

$$p_1 = p_0(1-\alpha_1)$$

After stage 2:

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

$$p_2 = p_1(1 - \alpha_2) = p_0(1 - \alpha_1)(1 - \alpha_2)$$

After K stages:

$$p_K = p_0 \prod_{k=1}^{K} (1 - \alpha_k)$$

4.2 Orchestration reliability metric (example)

Over N requests, define:

$$\text{Reliability} = 1 - \frac{\#\text{failed requests}}{N}$$

Further steps are required to build Autonomous Enterprises. The foundation to take the next steps forward consists of Digital Agents, Intelligent Agents, and Autonomous Agents. Research on risk evaluation and other Trust Factors, in particular for bias detection, and their automated supervision along with practical and concise validation and certification mechanisms, allows for a more scalable market growth of Digital Agents, Intelligent Agents, and Autonomous Agents in general. Efforts to integrate these earlier-autonomous primitives on a new orchestration layer, through simple APIs and specible Trigger Bus patterns associated with the individual interacting primitives, will increase Trust in these Engines, offer AI results that better fulfil user demand, and allow to Bijective integrate information and /BI-Proces refine and increase the resulting Train Set / Decision Quality Accuracy.

ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

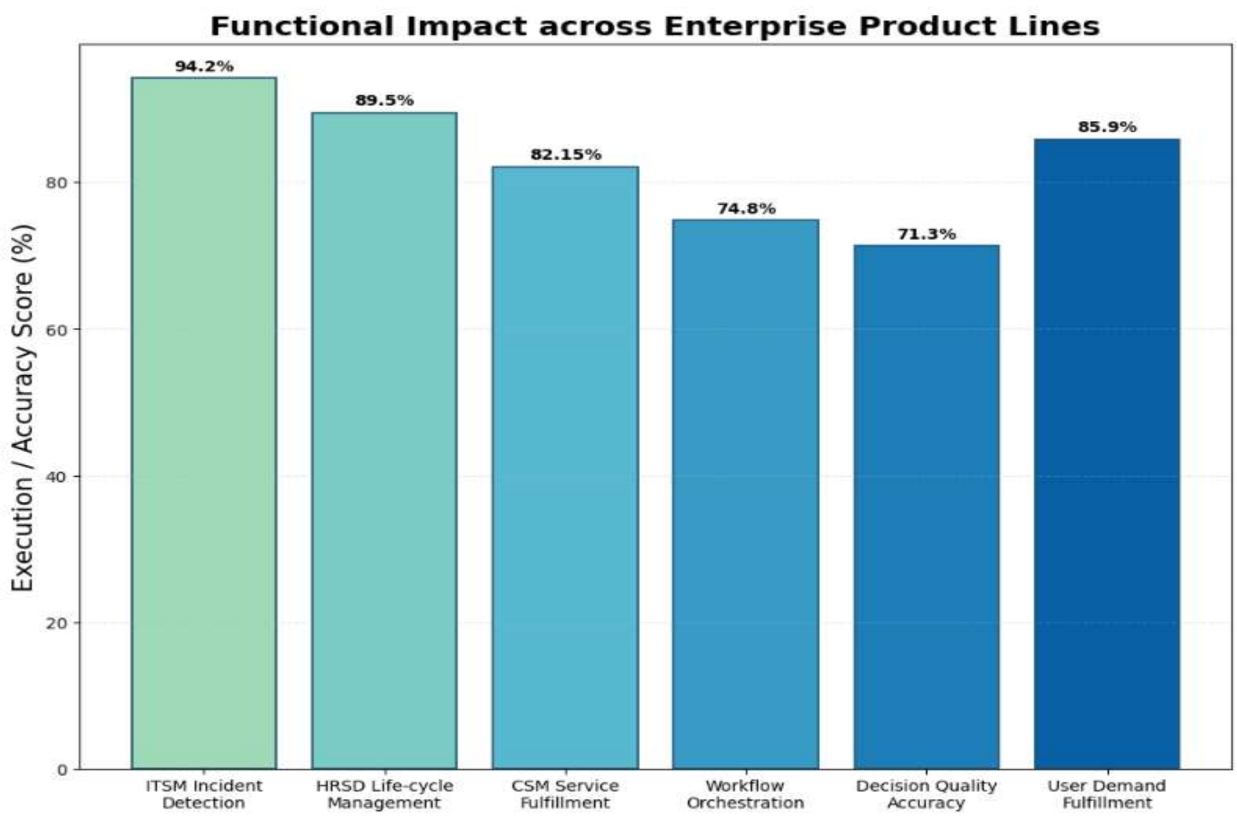

Fig 4: Functional Impact across Enterprise Product Lines

7.1. Key Takeaways and Future Directions

Group decision automation exposes several challenges. Reliability becomes critical to prevent cascading failures through the process. Transparent explanation of the decisions made can build user confidence and help identify weaknesses in the decision process. The objective is not to eliminate human discretion entirely but to avoid time and cost inefficiency resulting from incapacity or insufficient knowledge. Consequently, each automated decision should provide reasons for its logic, clear explanations of its recommendations, and opening for manual rollback procedure.

ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

The proposed approach is a first step into this direction. Current implementation focuses on operational deployment and performance improvement. Future work will involve extending the framework with new application domains, integrating data sources supporting the decision process, enhancing robustness of the decision evaluation, refinement, and confirmation, alleviating the reliance on orchestration in tracking the flow of data and execution, and strengthening support for incident management to allow automatic rollback of decisions.

References

1. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., ... Amodei, D. (2021). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877–1901.
2. 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. Advances in Neural Information Processing Systems, 35, 24824–24837.
3. Pamisetty, A., Paleti, S., Adusupalli, B., Singireddy, J., Inala, R., & Nagabhyru, K. C. (2025). Explainable AI Systems for Credit Scoring and Loan Risk Assessment in Digital Banking Platforms. In 2025 IEEE 13th International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS) (pp. 1478–1483). IEEE. 2025 IEEE 13th International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS). https://doi.org/10.1109/idaacs68557.2025.11322144
4. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. International Conference on Learning Representations.
5. Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2023). Toolformer: Language models can teach themselves to use tools. Advances in Neural Information Processing Systems, 36.
6. Chary, D. V., Meda, R., C, J. S. Mary., Narasimhachari, J. P., & A S, Y. (2025). TriFusionFormer: Tri-Modal Fusion Transformer Using Gated Modality Control and Multi-Scale Attention for Emotion Recognition. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1–8). IEEE. 2025

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

International Conference on Communication, Computer, and Information Technology (IC3IT). https://doi.org/10.1109/ic3it66137.2025.11341646
7. Mialon, G., Dessì, R., Lomeli, M., Nalmpantis, C., Pasunuru, R., Raileanu, R., Rozière, B., Schick, T., Dwivedi-Yu, J., Simões, C., de Almeida, M. B., & Scialom, T. (2023). Augmented language models: A survey. Transactions of the Association for Computational Linguistics, 11, 31–49
8. Alshar, M. M., Shahdadpuri, N., Rajeshwari, M., Gupta, M., Joshi, N. R., & Singireddy, J. (2025). Enhanced Management & Performance of Remote Workforce with Cloud and AI-Driven HR Analytics. In 2025 3rd International Conference on Advances in Computation, Communication and Information Technology (ICAICCIT) (pp. 631–636). IEEE. 2025 3rd International Conference on Advances in Computation, Communication and Information Technology (ICAICCIT). https://doi.org/10.1109/icaiccit68829.2025.11434104
9. Muthusamy, V., Rizk, Y., Kate, K., Venkateswaran, P., Isahagian, V., Gulati, A., & Dube, P. (2023). Towards large language model-based personal agents in the enterprise: Current trends and open problems. Findings of the Association for Computational Linguistics: EMNLP 2023, 6909–6921.
10. Li, X., Wang, S., Zeng, S., Wu, Y., Tang, Y., Wang, Y., & Yang, Y. (2024). A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges. Journal of Artificial Intelligence Research, 80, 1–50.
11. Krishnan, M., Nandan, B. P., Rongali, S. K., Meda, R., Kalisetty, S., & Singireddy, J. (2026, June). AI-Driven Data Engineering and Predictive Analytics Framework for Semiconductor Supply Chain Optimization and Digital Infrastructure Modernization. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1-6). IEEE.
12. 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. (2024). A survey on large language model based autonomous agents. Frontiers of Computer Science, 18, Article 186345.
13. Wang, F., Zhang, Z., Zhang, X., Wu, Z., Mo, T., Lu, Q., Wang, W., Li, R., Xu, J., Tang, X., He, Q., Ma, Y., Huang, M., & Wang, S. (2025). A comprehensive survey of small language models in the era of large language models: Techniques, enhancements, applications, collaboration with LLMs, and trustworthiness. ACM Transactions on Intelligent Systems and Technology, 16(6), Article 145.
14. Naik, A. V., Sheelam, G. K., Panchakatla, N., Muthukumaran, K., & Saranya, K. (2025). Comprehensive Analysis on Depression Detection From Social Media Using Deep Learning and Transformer Architectures. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1–8). IEEE. 2025 International

ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

Conference on Communication, Computer, and Information Technology (IC3IT). https://doi.org/10.1109/ic3it66137.2025.11341160
15. Nguyen, C. V., Shen, X., Aponte, R., Xia, Y., Basu, S., Hu, Z., Chen, J., Parmar, M., Kunapuli, S., Barrow, J., Wu, J., Singh, A., Wang, Y., Gu, J., Dernoncourt, F., Ahmed, N. K., Lipka, N., Zhang, R., Chen, X., ... Nguyen, T. H. (2025). A survey on small language models. Recent Advances in Natural Language Processing, 807–821.
16. Mukherjee, S., Mitra, A., Jawahar, G., Agarwal, S., Palangi, H., & Awadallah, A. (2023). Orca: Progressive learning from complex explanation traces of GPT-4. arXiv preprint arXiv:2306.02707.
17. Challa, K., Challa, S. R., Pamisetty, A., Kaulwar, P. K., & Koppolu, H. K. R. (2025, December). Transforming Payments: The Role of AI and Big Data in Fraud Alerts, Credit Monitoring, and Secure Transactions. In 2025 IEEE International Conference on Communication Networks and Computing (CNC) (pp. 1406-1414). IEEE.
18. Zhang, P., Zeng, G., Wang, T., & Lu, W. (2024). TinyLlama: An open-source small language model. arXiv preprint arXiv:2401.02385.
19. Ranaldi, L., & Freitas, A. (2024). Aligning large and small language models via chain-of-thought reasoning. Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics, 1812–1827.
20. Vadisetty, R., Nuka, S. T., Kalisetty, S., Pandugula, C., Burugulla, J. K. R., & Annapareddy, V. N. (2026). and Polypropylene Manufacturing. In Proceedings of Sixth International Ethical Hacking Conference: AI and Law (eHaCON 2025) (p. 269). Springer Nature.
21. Kim, S., Joo, S., Kim, D., Jang, J., Ye, S., Shin, J., & Seo, M. (2023). The CoT Collection: Improving zero-shot and few-shot learning of language models via chain-of-thought fine-tuning. Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, 12685–12708.
22. Fan, C., Tian, J., Li, Y., Chen, W., He, H., & Jin, Y. (2023). Chain-of-thought tuning: Masked language models can also think step by step in natural language understanding. Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, 14460–14473.
23. Maguluri, K. K. (2026). Cloud-Integrated Machine Learning System for Ebola Virus Disease Prediction and Epidemic Intelligence in Smart Healthcare Systems. Journal of Advances in Management, Engineering and Science (JAMES), 1(03), 1-9.
24. Cheng, X., Li, J., Zhao, W. X., & Wen, J. (2024). ChainLM: Empowering large language models with improved chain-of-thought prompting. Proceedings of the 15th International Conference on Language Resources and Evaluation, 1921–1931.

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

25. Yang, R., Song, L., Li, Y., Zhao, S., Ge, Y., Li, X., & Shan, Y. (2023). GPT4Tools: Teaching large language model to use tools via self-instruction. Advances in Neural Information Processing Systems, 36.
26. Rajamanickam, V., Singireddy, S., Davuluri, P. N., Sheelam, G. K., Aitha, A. R., & Vakkalagadda, T. (2026). AI-Enabled Visual Evidence Intelligence for Detecting Manipulated Digital Media. International Journal of Special Education, 41(14s), 115-124.
27. Haluptzok, P., Bowers, M., & Kalai, A. T. (2023). Language models can teach themselves to program better. International Conference on Learning Representations.
28. Wang, H., Ma, S., Dong, L., Huang, S., Wang, H., Ma, L., Yang, F., Wang, R., Wu, Y., & Wei, F. (2023). BitNet: Scaling 1-bit transformers for large language models. Proceedings of the International Conference on Learning Representations.
29. Frantar, E., Ashkboos, S., Hoefler, T., & Alistarh, D. (2023). GPTQ: Accurate post-training quantization for generative pre-trained transformers. International Conference on Learning Representations.
30. Davuluri, P. S. L. (2023). AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems (December 15, 2023).
31. Dettmers, T., Pagnoni, A., Holtzman, A., & Zettlemoyer, L. (2023). QLoRA: Efficient finetuning of quantized LLMs. Advances in Neural Information Processing Systems, 36.
32. Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., & Chen, W. (2022). LoRA: Low-rank adaptation of large language models. International Conference on Learning Representations.
33. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-T., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33, 9459–9474.
34. Bedi, B., Yandamuri, U. S., Kummari, D. N., Nagubandi, A. R., & Amistapuram, K. (2026, April). AI-Driven Sentiment and Behavior Analysis for Sustainable Business Growth. In 2026 International Conference on Emerging Research in Smart Electronics and Machine Informatics (ECMI) (pp. 1-11). IEEE.
35. Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Guo, Q., Wang, M., & Wang, H. (2023). Retrieval-augmented generation for large language models: A survey. arXiv preprint arXiv:2312.10997.

ISSN: 3067-4190

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 04 ISSUE: 02

RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

36. Asai, A., Wu, Z., Wang, Y., Sil, A., & Hajishirzi, H. (2024). Self-RAG: Learning to retrieve, generate, and critique through self-reflection. International Conference on Learning Representations.
37. Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Guo, Q., Wang, M., & Wang, H. (2024). Retrieval-augmented generation for large language models: A survey. ACM Computing Surveys.
38. Vankayalapati, R. K., Polineni, T. N. S., Ahammad, S. H., Pandugula, C., & Selvan, R. S. (2026). IoT-Enabled Augmented Reality for Real-Time Equipment Diagnosis. In Virtual Reality, Real Emergency (pp. 104-121). CRC Press.
39. 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. Advances in Neural Information Processing Systems, 36.
40. Shinn, N., Cassano, F., Gopinath, A., Narasimhan, K., & Yao, S. (2023). Reflexion: Language agents with verbal reinforcement learning. Advances in Neural Information Processing Systems, 36.
41. Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., Awadallah, A. H., & Wang, C. (2023). AutoGen: Enabling next-generation LLM applications via multi-agent conversation. arXiv preprint arXiv:2308.08155.
42. Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Wang, C., Zhang, Z., Wang, J., Yau, S. K. S., Lin, Z., Zhou, L., Ran, Y., Xiao, C., Wu, C., & Schmidhuber, J. (2024). MetaGPT: Meta programming for a multi-agent collaborative framework. International Conference on Learning Representations.
43. Qian, C., Liu, X., Wang, S., Li, H., Zhu, J., Liu, S., & Zhang, J. (2024). ChatDev: Communicative agents for software development. Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics, 151–167.
44. Schick, T., Dwivedi-Yu, J., Dessi, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., Cancedda, N., & Scialom, T. (2024). Tool use and tool learning in large language models. Foundations and Trends in Machine Learning, 17(4), 1–93.
45. Segireddy, A. R., Nagabhyru, K. C., Gadi, A. L., Pandiri, L., Paleti, S., Nandan, B. P., ... & Meda, R. (2026). U.S. Patent Application No. 19/389,116.
46. Qin, Y., Liang, S., Ye, Y., Zhu, K., Yan, L., Lu, Y., Lin, Y., Cong, X., Tang, X., Qian, B., & others. (2024). ToolLLM: Facilitating large language models to master 16000+ real-world APIs. International Conference on Learning Representations.
47. Patil, S. G., Zhang, T., Wang, X., & Gonzalez, J. E. (2023). Gorilla: Large language model connected with massive APIs. Advances in Neural Information Processing Systems, 36.

ISSN: 3067-4190
AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES
VOLUME: 04 ISSUE: 02
RECEIVED: APRIL 28 REVISED: MAY 13 ACCEPTED: MAY 21 PUBLISHED: JUNE 04

48. Li, M., Song, F., Yu, B., Yu, H., Li, Z., Huang, F., & Li, Y. (2023). API-Bank: A comprehensive benchmark for tool-augmented LLMs. Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, 3102–3116.
49. Deng, Y. (2023). IBM intelligent remediation for ITOps incidents powered by generative AI. NeurIPS 2023 Demonstration Track.
50. Agarwal, P., Dave, H., Bandlamudi, J., Sindhgatta, R., & Mukherjee, K. (2024). Multi-stage prompting for next best agent recommendations in adaptive workflows. Proceedings of the AAAI Conference on Artificial Intelligence, 38(21), 23189–23197.
51. Inala, R., Garapati, R. S., Aitha, A. R., Komaragiri, V. B., Gottimukkala, V. R. R., Recharla, M., ... & Varri, D. B. S. (2026). U.S. Patent Application No. 19/389,108.
52. Zhang, L., Jia, T., Jia, M., Wu, Y., Liu, A., Yang, Y., Wu, Z., Hu, X., Yu, P., & Li, Y. (2025). A survey of AIOps in the era of large language models. ACM Computing Surveys, 58(2), Article 44.
53. Wadhwa, A., Kaur, A., & others. (2024). Large language models for IT operations and AIOps: A systematic review. Journal of Systems and Software, 218, Article 112024.
54. Li, J., Li, G., Zhao, Y., Li, Y., & others. (2024). Large language models for software engineering: A systematic literature review. ACM Transactions on Software Engineering and Methodology, 33(5), Article 128.
55. Hou, X., Zhao, Y., Wang, S., Li, J., & others. (2024). Large language models for software engineering: A systematic literature review. ACM Transactions on Software Engineering and Methodology, 33(5), 1–58.
56. Wang, J., Wei, L., Wang, Z., & others. (2024). Large language models for business process management: Opportunities, challenges, and future directions. Business Process Management Journal, 30(6), 1701–1724.
57. Mendling, J., Decker, G., Hull, R., Reijers, H. A., & Weber, I. (2021). How do machine learning, robotic process automation, and artificial intelligence impact business process management? Communications of the Association for Information Systems, 48, 1–16.
58. van der Aalst, W. M. P. (2021). Process mining: A 360 degree overview. In W. M. P. van der Aalst (Ed.), Process mining handbook (pp. 3–34). Springer.
59. Zhang, L., Wang, S., Li, X., Chen, Y., & others. (2026). Large language model agents for autonomous enterprise workflows: Architectures, orchestration, and governance. Information Systems, 141, Article 102748.