



AgentOps for Governed Enterprise Service Delivery 1st Rama chandra nampalli , 2nd James Anderson

Rama Chandra Rao Nampalli, Solution Architect, nampalli.ramachandrarao.erp@gmail.com
James Anderson , Research Assistant

Abstract

Large language models (LLMs) are driving a rapid evolution in intelligent automation across industry sectors. Organizations are increasingly deploying LLMs to streamline operations by automating and augmenting human tasks and decision-making, with applications spanning the enterprise—particularly in service areas like IT service management (ITSM) and human resources service delivery (HRSD), where workflows tend to follow predictable, repetitive patterns. Yet automation becomes significantly more challenging in governed enterprise environments, where data privacy, security, and risk management are paramount; safety, compliance, and effective risk management must therefore be built into any advanced enterprise automation solution, not added as an afterthought.

Multi-agent LLM systems offer an effective framework for automating complex tasks while preserving this governance. Using a dedicated ITSM and HRSD task taxonomy, we show that multi-agent systems can automate a substantial share of incidents and service requests. Clearly defined agent roles and responsibilities help ensure data provenance, privacy, and security throughout the process: a central orchestrator manages end-to-end automation, while specialized agents handle dedicated functions such as safety and compliance monitoring for tasks involving sensitive data or materials.

keywords: Governed Enterprise AI; Intelligent Automation; IT Service Management; Human Resources Service Delivery; Safety; Compliance; Risk Management; Data Management; Large Language Models; Multi-agent Systems; ChatGPT; Negotiation; Orchestration; Task Taxonomy; ITSM; HRSD.

1. Introduction

In a governed enterprise AI environment, automation of IT service management (ITSM) and human resources service delivery (HRSD) offers potential for increased operational efficiency, cost-effectiveness, and improved reliability and user experience. The automation opportunities are significant since the ITSM and HRSD domains comprise repeated execution of well-defined tasks, often involving similar source data, frequently rely on external systems and interfaces, and are governed by a high degree of safety, compliance, security, and risk-management constraints.

The onboarding and offboarding of employees, together with access provisioning at the start and the termination of their employment, constitute a repetitive, high-volume HR operation that necessitates the continual handling of personally identifiable information (PII) and sensitive attributes. Task systems for these operations must therefore rigorously protect the data and respect



the underlying privacy legislation by enforcing an audit trail, retaining the data only for as long as necessary, and transferring it to secure environments for further archiving and modeling.

1.1. Onboarding, Offboarding, and Access Provisioning

Use of AI and automation for ITSM use cases such as user onboarding and offboarding, and provisioning and revocation of access for key applications.

Common tasks and activities in enterprise service environments typically follow established procedures with well-defined workflows and required documents. Such tasks are generally non-creative and repetitive, making them particularly suitable candidates for intelligent automation. For example, onboarding and offboarding users, and provisioning and revocation of access to applications that are subject to compliance controls, typically require repetitive interactions with multiple systems and respective business owners. These activities can involve sensitive data and have a direct impact on security and compliance, making intelligent automation in a governed manner especially important. Agent-based, multi-LLM automation of these processes provides appropriate levels of safety, control, and assurance, ensuring that the task can safely and successfully be completed.

Onboarding, offboarding, and granting or revoking access is generally performed in a de-centralized manner, resulting in time-consuming back-and-forth communications with employees and business function owners. For example, onboarding a new joiner typically requires the HR team to gather and consolidate several inputs from the onboarding manager (e.g., access to VPN, shared folders, and department-specific applications). The information then has to be communicated to the IT team to create the user account and provision access. Key accounts are created by the IT Security team, and the account in Active Directory is mapped to Azure Active Directory to enable single-sign on. Repeated checks with the respective application owners are required to ensure proper provisioning and delayed provisioning can impact productivity. Similarly, offboarding requires the HR team to regularly check and follow up with the respective application owners for required inputs to revoke access before the employee's last working day.

The provisioning and revocation of access for applications that require formal approval are also often times-consuming due to the necessary checks and follow-ups. For example, employee access to mainframe applications in financial services institutions is under strict audit scrutiny and provisioning requests typically require inputs from multiple senior managers. Business functions are required to create access matrix, specifying access to be granted to whom and the guiding principles for segregation-of-duties. The matrix is required for provisioning, as well as for subsequent forensic audits, to demonstrate adequate separation of duties including audit logs.




Fig 1: Multi-Agent LLM Systems for Intelligent

1.2. Employee Lifecycle and HRSD Automation

Automation of HRSD processes associated with the employee lifecycle is a primary use case for enterprise-AI-governed multi-agent technology. Human resource service-delivery actions required for onboarding, offboarding, and role changes—most importantly, access provisioning and deprovisioning—are frequent, operationally critical, and governed processes. These employee-lifecycle activities frequently generate challenges. New joiners and role changers expect timely access to the systems, platforms, and tools necessary to operate effectively. These requests need to be addressed in accordance with policy and security requirements. Delays in access provisioning often impede employee productivity, and flaws in denial of access can expose the organization to risk. In adequately governed enterprise-AI sectors, when timely, accurate, and secure completion of employee-lifecycle activities minimizes delay and risk exposure, full automation is a viable objective.

Additional employee-lifecycle changes are also candidate processes. Automated management of some of them can be a longer-term goal, with the greatest potential for improvement lying in the intersections with additional departments and systems. Extra touchpoints exist with travel offices, facilities management, and internal auditors, together with the full life-cycle record-



management conclusion workflow. The particular nature of deprovisioning requests means that employee-offboarding processes could be addressed by intelligent-agent techniques as well.

Equation 1: Operational Effectiveness

Let:

- C = task completion rate
- Q = quality score
- D = delay penalty
- all scaled between 0 and 1

We want effectiveness to:

- increase when completion increases,
- increase when quality increases,
- decrease when delay increases.

So begin with a weighted additive model:

$$E = w_1C + w_2Q - w_3D$$

where $w_1, w_2, w_3 \geq 0$.

Now normalize the weights so the score stays interpretable:

$$w_1 + w_2 + w_3 = 1$$

Then the final short equation becomes:

$$E = w_1C + w_2Q - w_3D$$

Step-by-step logic

Start with the verbal definition:

$$\text{Effectiveness} \propto \text{completion} + \text{quality} - \text{delay}$$

Replace words by symbols:

$$E \propto C + Q - D$$



Add controllable importance weights:

$$E = w_1C + w_2Q - w_3D$$

2. Background and Foundations

Enterprise AI environments are governed by corporate policies that seek to minimize risk. Two critical business processes that require efficient support but are often under-prioritized are IT service management (ITSM) and human resources service delivery (HRSD). These processes are governed by well-established frameworks and policies but can vary greatly in detail, mainly focusing on incident and request management. A task taxonomy for ITSM and HRSD encompasses the full employee lifecycle, including onboarding and training. User experience can be enhanced if data protection regulations are respected and using automated agents is optional for employees.

Dedicated enterprise service desks specializing in IT service management and human resources service delivery ensure efficient, specialized, and reliable handling of business-critical requests. However, despite the multifaceted nature of business processes, such specialized service desks are commonly governed by toolsets optimized for simple, single-person, non-sensitive tasks. Operational effectiveness and reliability are just two essential metrics for a service desk in a regulated enterprise; compliance with laws, regulations, and internal policies is paramount. Sensitive usage data must be stored and processed responsibly as well. Although data provenance—clear and specific recording of data sources and quality—is equally vital, it is seldom integrated into enterprise service desks.

2.1. Enterprise AI Governance

Enterprise AI requires governance to mitigate risks arising from AI misuse, mitigate governance and cyber-security risks, protect sensitive training data, and ensure compliance with regulations governing the enterprise domain. Board-Driven AI Governance is critical for a successful Enterprise AI strategy and is becoming standard practice within governed business environments. Enterprise AI Governance must manage data lineage, privacy, and data security and impact the entire dataset lifecycle, external data usage, pre-training checks, prompt design, response validation, and core and ancillary dataset updating.

Governance requirements strongly influence architecture design for Enterprise AI, particularly in regulated domains such as finance, insurance, pharmaceuticals, and government. The presence of sensitive data requiring anonymization or masking before LLM training, deployment, or usage extends beyond the training dataset and application-specific data. It requires specific handling within all LLM workloads, especially those integrating training, validation, and application of data extracted.




Fig 2: Managing Multi-Agent LLM Systems in Enterprises

2.2. IT Service Management and Human Resources Service Delivery

IT service management (ITSM) comprises processes and practices that cover the end-to-end delivery of information technologies to the business customer. ITSM emphasizes the delivery of tech support services and service requests. Incident management helps to restore a service and reduce adverse business impacts. Request management fulfills user service requests—standardized requests for information, advice, or a minor change. Change management ensures that any changes to the business's IT environment or services are controlled, coordinated, and carried out with minimum risk. Problem management reduces the risk of recurring incidents. The core processes of ITSM, along with supporting resources used to ensure their effectiveness and efficiency, form a service management system.

Human Resource Service Delivery (HRSD) describes the many services that HR provides to the organization, and its associated processes, systems, and tools. It is a continuum of service delivery that includes everything from a single hot line for help and answering questions, through a service center design that provides help along with some fulfillment (for example, change of address), through a service hub that provides help, more-broad-based fulfillment services, and gives an employee-video, payroll, and time-information services. The concept of HRSD encompasses everything that HR does in regard to delivering services to employees.

Equation 2: Operational Reliability

Let:

- P_s = probability of successful execution



- P_t = probability of meeting turnaround/SLA time
- e = error rate

Since reliability should drop when error increases, use $1 - e$.

Assuming these three factors jointly determine reliability:

$$R = P_s \cdot P_t \cdot (1 - e)$$

Step-by-step derivation

Reliability depends on three simultaneous events:

- success,
- timeliness,
- low error.

So:

$$R \propto (\text{success})(\text{timeliness})(\text{low error})$$

Write low error as:

$$\text{low error} = 1 - e$$

Substitute symbolic terms:

$$R \propto P_s P_t (1 - e)$$

Use direct normalization:

$$R = P_s P_t (1 - e)$$

3. System Design Principles for Governed Environments

Enterprise ITSM and HRSD services often operate under strict compliance and governance regimes, so the systematic application of agent-based LLM technology requires careful consideration of safety, compliance, and risk management. Specific, domain-expertising and insurance-type prompt layers are also necessary to make the LLMs controllable in such environments and to facilitate the dynamically negotiated anonymisation of input data for all agents supporting ITSM and HRSD functions.

Safety, Compliance, and Risk Management Enterprise ITSM services are often treated as part of system and network operations, which includes the management of incidents, changes, problems, configurations, changes, and SLAs. Users request services, and operators fulfil requests, which have no external customers. Responsibility for authorisation of changes and decisions is normally borne by change management teams, but the human checks and balances—and consequently the insurance—can be relaxed when automating low-risk changes when augmented by other ITSM functions. User requests for service can impact other users and the integrity of IT services, but with appropriate training and insulation of the LLM prompt it is reasonable to treat LLMs as a specialised low-collision incident/malicious activity detection and prevention service.

However, this assumption does not hold for HRSD services, which deal directly with employee data; any breach could have severe consequences for the users, their employers, and the company providing HRSD services. Security and compliance are paramount, requiring a broker mechanism that tightly controls the flow of sensitive data into the HRSD services. The presence of Sensitive Data By Country in ITSM also introduces risks, as ITSM functions such as diagnostics and service request responses can leak sensitive information in certain contexts. To maintain compliance without auditing every user transaction, these sensitivity requirements can be dynamically applied not by an LLM pre-definition during training but by the normal human process of modifying the contract for an federated agent that the user is negotiating with, introducing temporary Data Privacy Risk Conditions that require anonymisation of all data feeds being passed to the agent.




Fig 3: System Design Principles

3.1. Safety, Compliance, and Risk Management

Governed enterprise IT environments mandate automated safety, compliance, and operational risk management for deployed ITSM and HRSD systems. This requirement prevents automated agent interactions from establishing any risk vectors not



previously present, avoiding new hazards arising from the AI system's integration. Storage and processing of sensitive data—notably identity information—bring additional risk, and specific steps support safe automation in these cases.

Agent-requested information for incident resolution cannot include personally identifiable user details, and any data set at risk or subject to particular process conditions needs special attention. A dedicated data-filtering agent verifies that safety concerns have been mitigated, checking three operating-principle conditions over the agent interactions. Firstly, data access should not introduce race conditions. Secondly, data access for knowledge-gathering must occur after other agents have acknowledged their requests. Finally, if any agents have placed specific risk conditions on the incident, those conditions must be satisfied before knowledge discovery.

Category	ITSM (IT Service Management)	HRSD (Human Resource Service Delivery)
Core Focus	IT services delivery & support	Employee lifecycle services
Key Processes	Incident, Request, Change, Problem	Onboarding, Offboarding, Access provisioning
Data Type	System/service data	Highly sensitive PII data
Risk Level	Medium–High	Very High
Automation Suitability	High (structured workflows)	High but governance-heavy
Compliance Needs	Moderate (security + logs)	Strict (GDPR, privacy laws)

Table 1: ITSM vs HRSD Task Classification Table

3.2. Data Lineage, Privacy, and Security

Enterprise ITSM and HRSD operations involve the use of sensitive personal and organizational data, which are subject to various regulations that dictate how they need to be treated. For example, the GDPR requires that personal data be kept private and secure, and it imposes strict rules on whether, when, and how data can be stored, shared, and used. The entire lifecycle of the data must be recorded so that a data controller or processor can present an auditable trail upon request. The same is true for data covered by HIPAA and other regulations. Therefore, data traceability and secure handling are crucial.

Managing data privacy, security, and lineage requires functionality in the agents performing specialized tasks. The data sources and attributes required for ITSM and HRSD, as well as privacy and compliance constraints, are part of the specification of ITSM and HRSD tasks. Security measures for data storage and transmission address data privacy and security.

Equation 3: Compliance-Governed Performance

Step-by-step derivation

Start from the paper’s logic:

$$\text{Governed performance} \propto \text{effectiveness} \times \text{reliability} \times \text{compliance}$$



Replace with variables:

$$G \propto E \cdot R \cdot C_p$$

Use direct normalized form:

$$G = ERC_p$$

4. Task Taxonomy for ITSM and HRSD

Any enterprise with a sizable IT department or human resources organization regularly deals with large volumes of repetitive requests for incident resolution, service provisioning, or information retrieval. Although automation of these tasks is of considerable interest, service delivery automation remains limited in both the ITSM and HRSD areas. The root cause can be found in a mismatch between the underlying technology and the requirements of such tasks: they require not only a knowledge base but also an understanding of natural language and negotiation skills, qualities that make them well suited for Large Language Models.

Automation of HRSD is especially challenging due to the need to operate on typically sensitive data surrounding employee onboarding and offboarding, provision of access, and retrieval of payslip or tax-related information. Products are beginning to emerge for the secure use of LLMs in enterprise contexts, but there is yet no clear solution for the specific automation of HRSD tasks. Such automation also typically needs to meet the stringent governance requirements of a financial services institution, including regulations around data validation, auditability, consent, and illicit data processing.

4.1. Incident and Request Management

An incident is an unplanned interruption to a service or a reduction in the quality of a service, while a service request is a user-initiated action that results in the provisioning of new or changed access to IT services or products. The fundamental difference between the two is the underlying impact on the enterprise services. For example, a request for a new user account represents service change creation, potential change fulfilment, and privacy/sensitive data handling requirements. Conversely, a user account deletion is an incident, threatening enterprise service availability and continuity unless the request is for an already scheduled user onboarding situation.

Achieving operational excellence requires pre-emptive mitigation of incidents combined with hassle-free fulfilment of requests. The complete resolution of any major business-affecting incident can take a considerable amount of time, while a user whose password has expired should not have to wait a day for access to the services. Often, service restoration can be speeded up for major incidents by completing selected steps in parallel. Such parallelisation requires that the possible steps be assigned to person(s) capable of completing the step(s).




Fig 4: Incident and Request Management

4.2. Change and Problem Management

Fundamentally, significant problems experienced by any enterprise imply a lack of stability that can result in substantial impacts on IT services and consequently on the enterprise's ability to conduct business. Change is the ability to constantly introduce new services and modify the existing ones to satisfy the present and future requirements of the business at optimal resource cost. Change and Problem Management processes correctly designed and applied allow a business to protect the service quality and support the strategy of change by keeping the risks linked to modification at a manageable level. Problem Management determines the root cause of reusable incidents and minimizes the impact of incidents with known errors on the business. Various stages in the process cycle are often operated with heterogeneous solutions different for each stage. Multiple solutions in a distributed way for any task of the management of these processes could increase efficiency and reduce costs. However, these solutions must be used in a coordinated way to achieve the objectives of the processes. A multi-agent approach to the coordination and orchestration of AI solutions for Change Management and Problem Management thus increases the ability mechanisms behind these processes. An operational architecture embedding these concepts is sketched.

The stability of IT services is a prerequisite for enabling business continuity, while their visibility is crucial for the final users' experience. Conversely, substantial volatility implies that some processes must be monitored to warn users about possible failures, including the failure of any Service element. Major Incidents are thus recognized in the operation of a service that have a significant impact on the business, lead to suspension of a Service or a reduction of Service Quality, require the involvement of several technical support groups and are not resolvable within a time period agreed with the business. Service Management processes efficiently working with a modified operational model are thus necessary to keep the incidents attractive for customers and to minimize their impact on business. Nevertheless, those processes must also guarantee the ability to provide new or modified Services and to maintain known errors at an acceptable level.

Equation 4: Privacy-Risk Reduction

Let:

- P_0 = baseline privacy risk before controls



- a = anonymization effectiveness
- f = filtering effectiveness
- b = broker/control effectiveness

Each control reduces remaining risk. So after anonymization:

$$P_1 = P_0(1 - a)$$

After filtering:

$$P_2 = P_1(1 - f)$$

Substitute P_1 :

$$P_2 = P_0(1 - a)(1 - f)$$

After broker control:

$$P_r = P_2(1 - b)$$

Substitute again:

$$P_r = P_0(1 - a)(1 - f)(1 - b)$$

So the final equation is:

$$P_r = P_0(1 - a)(1 - f)(1 - b)$$

5. Agent Coordination and Orchestration

Centralized orchestration for multi-agent systems can simplify the design of negotiation protocols, coordination strategies, and mechanisms for ensuring the reliability of multi-agent systems. A different style emerges when implementing a lightweight, decentralized coordination mechanism among the agents. Relatively simple techniques can implement intra-agent coordination and establish negotiation and delegation policies.

Multi-agent systems are often designed to work in a centralized way, where an orchestrator coordinates the actions of each agent. A central orchestrator helps in the development of multi-agent systems through the difficulties introduced during the definition of coordination between agents. However, a detailed perspective considering a distributed approach provides alternative insights. The integration of a proposal with a central orchestrator and a decentralized coordination mechanism can improve some

operational aspects of the multi-agent system. A hierarchy of tasks can originate the need for negotiation, delegation, and the required rules to achieve such goals. A key aspect in such systems is how to make it operationally effective, efficient, and reliable.




Fig 5: Agent Coordination and Orchestration

5.1. Central Orchestrator versus Distributed Coordination

Agent coordination and orchestration can be implemented either by a central orchestrator or via distributed approaches. Central orchestrators focus on the complete automation of processes, while distributed systems enhance a collaborative environment where agents negotiate, share information and context, and delegate tasks. Multi-agent LLM-based systems evolve toward the distributed coordination of communication agents focused on incident and request management, as well as on informal change requests. Central orchestration remains an efficient solution for contact-center-style interaction with users and for fully automated processes defined by the enterprise policies.

A central orchestrator can be beneficial for the following use cases: scenarios defined by the enterprise policies and involving sensitive data, where the request and answers strictly follow the transactional communication model; scenarios involving users with limited skills for business-process navigation, such as end users requesting expert-access provisioning or business leads (i.e., change manager, procurement) reporting issues with the fulfillment of changes or requests. For incident and request management, the standard transactional structure is typically preferred; however, interaction often evolves toward a dialog, making a communication-agent approach feasible.

Phase	Tasks	Systems Involved	Risks
Onboarding	Account creation, access provisioning	HR, IT, Security, AD/Azure	AD Data leakage, delay
Role Change	Modify access rights	HR, IT, Application owners	Incorrect privileges



Phase	Tasks	Systems Involved	Risks
Offboarding	Revoke access, deactivate accounts	HR, IT, Security	Security breaches
Post-exit	Audit logs, data archiving	Compliance systems	Regulatory violations

Table 2: Employee Lifecycle Automation Tasks

5.2. Negotiation, Delegation, and Conflict Resolution

Depending on the task species, distributed MAS, a potentially richer alternative than centralized design, may better serve safety and data compliance objectives. Yet it raises the need for negotiation, delegation, and conflict resolution mechanisms that support the agile adaptation and real-time operation required by service management workflows, while ensuring these remain safe and compliant in governed enterprise environments.

Central orchestration has enabled real-time automation of ITIL incident and request operations in production and, for other ITIL workflows, user-defined priority-based regionally-centric automatic negotiation between local instances of multi-agent systems. Two different MAS have been developed for ITIL management of data-sensitive employee onboarding-offboarding-access provisioning operations: a distributed MAS supporting simultaneous real-time operation with a demilitarized zone agent for compliance purposes, and a central MAS that foresees delegation to external cautious agents. These implement an incident request duo servicing a safe, minimal-local-data-life-cycle policy for formal enterprise compliance purposes, and are extendable to related operations such as IT and security policy development and deployment. The global querying and backup and maintaining of critical digital assets work for such sensitive tasks.

Equation 5: Safety Gate for Agent Knowledge Access

Let:

- $N_r = 1$ if there is **no race condition**, else 0
- $A = 1$ if acknowledgements are complete, else 0
- $K = 1$ if risk conditions are satisfied, else 0

Then access is allowed only when all three are true:

$$S = N_r \wedge A \wedge K$$

If we want the same idea in algebraic 0/1 form:

$$S = N_r \cdot A \cdot K$$

Step-by-step derivation



Condition 1:

$$N_r \in \{0,1\}$$

Condition 2:

$$A \in \{0,1\}$$

Condition 3:

$$K \in \{0,1\}$$

All must hold together, so use logical AND:

$$S = N_r \wedge A \wedge K$$

For binary variables, AND is multiplication:

$$S = N_r A K$$

Hence:

$$\boxed{S = N_r A K}$$

6. Data Management in Governed Environments

Effective data management addresses both input data to the agents and output data generated by the agents. The quality of foundation models can be sensitive to the datasets used for fine-tuning. In the case of the ITSM and HRSD tasks, while foundation models are not directly tuned, the quality of the generated responses depends heavily on training, test, and validation datasets and the similarity between the distribution of incoming requests and the original data distribution. These distributions are driven by the data sources feeding the requests into the agents. For ITSM and HRSD capabilities actively running in production use, tampering sources lead to increased error rates and negatively affect the user experience. In heavily regulated enterprise environments, the data generated by the tasks and the logs produced need special attention. For some regulatory frameworks, these logs contain sensitive data and must be treated according to the respective compliance and data protection laws.

6.1. Data Sources and Quality for ITSM and HRSD

Several aspects should be monitored in terms of data quality for these automation areas: requests coming from different sources must be monitored and optimized; if specific sources present a different distribution from the normal one, more care should be taken; if there is an incoming request that never comes from training, a special focus should be put on the response; regular domain-oriented reviews of the requests and respective responses must be performed; and all training test lists must be open for



contributions from external stakeholders. The monitoring of these aspects allows the detection of data-shift problems, of areas that are more critical than others, and of alerts that should be generated due to a lack of fine-tuning care of segments with less data.

6.2. Data Privacy, Anonymization, and Compliance Requirements

Several compliance and data protection laws and regulations require that personally identifiable information, sensitive information, and categories of information that the organization desires not to disclose be redacted from logs and database records. The logs that fall under these rules must be kept private and protected. Preventive measures for generating data that contains personally identifiable information, proprietary information, confidential information, or other sensitive information must be in place and help the business reduce risk.

6.1. Data Sources and Quality for ITSM and HRSD

Incident management relies chiefly on internal systems and databases. A limited language model suffices for high-availability use cases, while a broader language model can process incident details, directs the LLM-based agent to related knowledge base articles, and check that the response is both factual and helpful, thereby supporting knowledge management processes. Request management additionally requires a model of offered products, services, SLA expectations, and user entitlements within the service catalogue. Change management and problem management processes make extensive use of historical incident and request data. Routine change requests can be automated based on models or classifiers, and an LLM-based agent can evaluate both patterns of recurring incidents and known errors according to change and release management processes. Security-related problem analysis requires integration of sensitive data sources, such as malicious domains, phishing URLs, malware categorization of files, and recent attacks. A centralized data platform, governed testing and validation processes, and sound generative and recognition prompt design improve data quality and test process exhaustiveness.

HR service delivery automation, particularly onboarding, offboarding, and access provisioning, cannot ignore data privacy. Though the SSO transition minimizes sensitive data leakage, onboarding requires the passage of birth dates and gender in public requests, access provisioning must support roles containing hunting-dog words, and offboarding necessarily entails the leaking of personally identifiable information. A system built for such sensitive domains must guarantee data protection at all stages: prompt design, fallback option selection, training and input data collection, output review, and detection of future user satisfaction or dissatisfaction. The privacy-preserving data-sharing blockchain framework provides one approach to protecting sensitive data shared by untrusted participants in decentralized environments. A stronger verification mechanism permits more flexibility in accessing only certain sensitive data within the shared pool.




Fig 6: Data Management in Governed Environments

6.2. Data Privacy, Anonymization, and Compliance Requirements

Anonymization of sensitive data is often critical for meeting enterprise privacy, security, and compliance requirements. While potentially useful for training or fine-tuning generative LLMs, such techniques can ultimately hinder the fidelity of the data and models, unless additional care is taken in modeling and generating their outputs. The adoption of more complete LLM-based multi-agent solutions in the enterprise, especially in highly governed domains like ITSM and HRSD, may nevertheless enable processes that leverage raw, sensitive data, while addressing privacy, security, and compliance explicitly—achieving the opposite of traditional anonymization. LLMs also support different techniques for redacting sensitive data in a more targeted, fine-grained fashion. Such data leakage controls can be defined for end-users or external third parties, and can make use of only LLM



capabilities for unstructured data, physical isolation (i.e., separation of the content) for structured data, and well-trained, privileged agents for additional sources.

Two obvious examples for the application of such techniques are solutions for HRSD and ITSM request-response discussions. In the case of HRSD, redundant information about employees being offboarded is disseminated across different agents—certain HR oriented ones generating the language content for requests (and responses) and others—beyond the obvious provisioning and de-provisioning ones—checking the free-text content for sensitive personal content, or just for semantically-scanning or re-authoring regulatory-document-level responses. Similarly, for ITSM incidents, these techniques are instrumental for incident discussion threads: generative models for incident detection and for natural-language discussion both reside with the ITSM broker, while detection of sudden leaks of business-sensitive data is reallocated to a separate agent outside the ITSM environment (and potentially also outside the enterprise).

Flow chart :

[Start]

|

v

[Employee lifecycle event occurs]

|----> New hire

|----> Role change

|----> Exit / termination

|

v

[HRSD request is initiated]

|

v

[Central Orchestrator receives request]

|

v



[Identify request type]

|----> Onboarding

|----> Access provisioning

|----> Access revocation / offboarding

|

v

[Collect required inputs]

|----> Employee identity data

|----> Manager inputs

|----> Department/application needs

|----> Compliance requirements

|

v

[Safety / Compliance / Privacy check]

|

+---- No ----> [Request missing or unsafe]

| |

| v

| [Ask for corrections / approvals]

| |

| v

| [Re-check]

AMERICAN DATA SCIENCE JOURNAL FOR ADVANCED COMPUTATIONS

VOLUME: 01 ISSUE: 01

RECEIVED: OCTOBER 14

REVISED: OCTOBER 28

ACCEPTED: NOVEMBER 02

PUBLISHED: NOVEMBER 18



|

+---- Yes ----> [Proceed]

|

v

[Route work to specialist agents]

|

+-----+-----+-----+

|

|

|

v

v

v

[HR Agent]

[ITSM Agent]

[Security/Compliance Agent]

|

|

|

|

|

|

+-----+-----+-----+

|

v

[Create / update user identity]

|

v

[Provision apps, VPN, folders, shared resources]

|

v

[Map identity to SSO / directory systems]

AMERICAN DATA SCIENCE JOURNAL FOR ADVANCED COMPUTATIONS

VOLUME: 01 ISSUE: 01

RECEIVED: OCTOBER 14

REVISED: OCTOBER 28

ACCEPTED: NOVEMBER 02

PUBLISHED: NOVEMBER 18



|

v

[Request app-owner confirmations]

|

v

[Check segregation-of-duties / approvals]

|

+-----+-----+

| |

No Yes

| |

v v

[Escalate to human owner] [Execute provisioning]

| |

+-----+-----+

|

v

[Write audit trail / lineage logs]

|

v

[Notify employee, manager, HR, IT, security owners]

|



v

[End]

7. Evaluation Metrics and Validation

Effectiveness and reliability represent central requirements for Multi-Agent LLM (MALLM) systems. Operational effectiveness corresponds to a system's capacity to deliver expected outcomes for users, while operational reliability describes a system's ability to achieve these outcomes consistently over time. Non-invasive, service-based MALLM systems fulfill many IT-Service Management (ITSM) and Human Resources-Service Delivery (HRSD) tasks without complex training or setup of LLM components. Success thus hinges on individual agents' competence and availability for task execution.

Safety and compliance represent further critical dimensions of operational validation for MALLM systems deployed in governed environments. For many user-facing tasks, such as incident management (with ITSM) or employee onboarding (with HRSD), applying best-practice procedures represents a baseline requirement. Achieving acceptable outcomes for sensitive tasks involving Private Identifiable Information (PII) or human-like information generation calls for additional control mechanisms. For simple inquiries, User Data and Model Use protocols allow non-sensitive data requests to explicit user requests. For PII or user data-related questions, user description of security, privacy, and data governance rules should accompany the user prompt. To validate output compliance, an automatic data-governance policy must also generate domain-focused requirements.

7.1. Operational Effectiveness and Reliability

Multi-agent systems are typically evaluated in terms of the desired capabilities and performance indicators. However, LLM agent systems also require specialized evaluation dimensions. That is the case for ITSM and HRSD automation in enterprise AI applications governed by data privacy, security, and regulatory compliance requirements.

ITSM and HRSD service deliveries at operational levels must be executed effectively and reliably. Effectiveness covers the ability to fulfil the end-to-end request, assuring the desired level of quality and in the expected time frame. Reliability encompasses the likelihood that a task is successfully executed within the acceptable turnaround time and error rates. Therefore, a suitable combination of these aspects reflects operational effectiveness and reliability. Here, effectiveness quantifies the successful completion of tasks and the observed delays, identifying the factors affecting the results. Reliability measurements focus on the successful execution of operations in a timely manner and at a low error rate, drawing attention to the main stress points. Operating under the aforementioned compliance constraints adds further pressure, with user data privacy, internal data security, or regulatory adherence possibly causing explicitly detectable IM failures.


7.2. Compliance and Governance Metrics

Compliance with enterprise policies and standards is vital for mitigating financial, operational, and reputational risks, safeguarding customer confidence, and enabling digital transformation. Multi-agent systems should thus ensure compliance with corporate policies, address contractual and regulatory obligations, and account for Change Management, Problem Management, Incident Management, and Access Provisioning policies specified by the relevant ITIL frameworks. Supervision—conceptualized as the extent of monitoring or oversight necessary to ensure compliance with company policies and other rules—is essential, as are explicit trust relationships between agents and external entities. A deliberate account of the data lineage for all tasks also

informs compliance considerations. Naturally, any data used for training or testing agent models must conform to data protection standards.

Beyond these safety and compliance concerns, enterprise AI solutions need to provide governance metrics that assemble and quantify auditing information from across the entire agent ecosystem. Each agent should supply distinct auditing logs related to their expertise and activity, which in turn enable stakeholders—including compliance, risk management, information security, and data privacy teams—to build a consolidated governance view specific to their remit or responsibility.

AI Performance Dashboard - Enterprise View



7.3. User Acceptance and Experience

User acceptance of intelligent automation solutions directly affects their success and, importantly, constitutes an evaluative



measure in its own right. User acceptance depends not only on the operational effectiveness and reliability of the integrated agent systems but, from the other side, also on the degree to which the systems enrich user experience.

The automation of incident and request management processes using the LLM-based agent systems was implemented in a regulated enterprise operating within the pharmaceutical sector. Supporting agent cooperation took place via a central orchestrator capable of real-time assessment of all interactions. Feedback from end users indicated user satisfaction with resolution quality and speed. Furthermore, interaction and resolution accuracy levels were sufficiently high that some users opted not to engage with the IT support team.

8. Practical Implementation Scenarios

Using the aforementioned principles, a taxonomy of interactions among ITSM and HRSD agents has been created to automate tasks within a governed enterprise environment. Several practical implementation scenarios support the proposed taxonomy, each of which reuses existing ChatGPT open-source agent implementations²⁸ and combines them in new ways to perform enterprise tasks according to the defined interaction models.

The first scenario focuses on Hosting Provider IT Service Management. Here, a central orchestrating agent deploys ChatGPT ITSM agents specialized for the Service Desk and Change Management functions. These agents communicate directly with external-organization agents (ChatGPT agents set up by users in a different ChatGPT organization) to validate user identity for incidents and requests, return user-sensitive data, and obtain system- or user-sensitive data that can be obtained from other sources outside the hosting provider. The ability to check the OpenAI Data Usage Policy and manage the data allowed for processing helps address crucial safety and compliance issues when automating processes that handle highly sensitive user data. Data management requirements are based on the OpenAI Privacy Policy.

Flow chart:

[Board / Enterprise Governance]

|

v

[AI governance policies]

|----> Privacy

|----> Security

|----> Data lineage

|----> Ethics



|----> Compliance

|

v

[Enterprise AI architecture constraints]

|

v

[Central Orchestrator]

|

+-----+-----+-----+

| | | |

v v v v

[ITSM Agents] [HRSD Agents] [Compliance Agents] [Data Agents]

| | | |

+-----+-----+-----+

|

v

[Enterprise systems]

|

+-----+-----+-----+

| | | |

v v v v

[HR systems] [Service desk] [Identity stores] [Knowledge bases]



|

v

[Responses / actions]

|

v

[Monitoring, audit, validation, logging]

|

v

[Governance review and policy feedback loop]

|

v

[End]

8.1. Case Study: ITSM Automation in a Regulated Enterprise

A case study illustrates the application of a multi-agent system for automating ITSM incident and request responses in a large European bank. The technical architecture employs multiple LLMs. Core incident and request management functions are implemented in a central orchestration agent that also serves as a user-facing application. All connections to the service management tool are performed via this agent. Several supporting agents assist with quality assurance during response generation, validate incident data for acceptance, and perform other activities such as checking for related incidents and problems, or resolving detected problems via a Problem Management Change Process.

The bank operates under strict regulatory controls that include data residency constraints disallowing the transfer of production environment data outside Switzerland. To meet the fact-checking requirement using LLMs, user prompts are routed to a Swiss-based instance of ChatGPT. Software for checking incident data format, quality, and acceptance has been built and tested. Other quality checks are currently being implemented and the first round of user acceptance testing is planned.

8.2. Case Study: HRSD Automation with Sensitive Data


A second practical implementation scenario considers the onboarding, offboarding, and access provisioning sub-processes of HRSD automation in a governed enterprise environment. The objective is to establish a coherent framework that supports the sensitive nature of the data involved. Compliance with data residency and privacy acts such as the EU GDPR mandates the establishment of effective data anonymization mechanisms that ensure full transparency. During the implementation of these



three HRSD sub-processes, the achieved anonymization enables the legitimate use of public Large Language Model (LLM) services, thus avoiding vendor lock-in and fulfilling both legal and executive compliance.

Enterprise HRSD systems govern many aspects of the entire employee life cycle, from recruitment through onboarding, access provisioning, and offboarding to eventual separation from the organization. Historically, intra-organizational reputation and social pressures ensured that access provisioning occurred seamlessly. Recent world events exacerbated by a talent shortage, however, have led to a growing number of requests for immediate HRSD fulfillment. Sourcing, recruiting, onboarding, supplying, training, retraining, and upskilling need to coincide with both formal LLM systems that provide guidance only through text and internal corporate values.

Enterprise AI Data Usage in 2025



9. Conclusion

The value of enterprise AI has received significant attention in recent years. Early applications focused on cost reduction, productivity improvement, and employee experience. However, there is now a greater emphasis on using enterprise AI to provide



advice and guidance in complex, regulated domains. Practical examples include banking, insurance, finance, and health care, where the risk associated with inappropriate AI use can have serious ramifications.

With many organizations in these markets undergoing digital transformation, new designs are needed to deploy AI LLM-based systems in accordance with corporate governance requirements. For these domains, governance-related aspects must be actively managed, especially in the areas of safety, compliance, and risk management; data lineage and privacy; and auditability and traceability. Such designs can be enhanced by applying a multi-agent systems perspective to tasks like ITSM and HRSD automation, where safety, compliance, and privacy-related aspects can be effectively addressed by supporting systems and agents.

9.1. Emerging Trends

The growing societal awareness of the risks and limitations associated with generative AI, with calls for regulations and governance, is leading enterprises to implement controlled and risk-aware enterprise-wide AI strategies. These strategies aim to leverage the benefits of artificial intelligence while minimizing data exposure and ensuring content compliance. Consequently, many organizations have focused on establishing an enterprise AI governance framework based on formalized risk assessment and management approaches. The growing concerns over the quality of the content produced by large language models have further heightened the need for such governance approaches in generative AI scenarios.

One possible avenue for implementing enterprise AI governance is to support the orchestration of a wide range of tasks by leveraging multi-agent systems that combine multiple large language models. These systems simplify and facilitate the control over the interaction with the models, as well as the incorporation of external resources needed for task completion. However, enterprise AI governance imposes additional requirements on both the content-generation phase and input data privacy. In addition, enterprise generative AI frameworks are often designed only for specific use cases, while enterprise environments typically support interaction with end-users in the context of human resources service delivery and information technology service management processes. Therefore, taxonomies of these interactions and the associated management, decision-making, and orchestration challenges must be defined to facilitate their automation.

10. References

1. Chen, Z., Kang, Y., Gao, F., Yang, L., Sun, J., Xu, Z., Zhao, P., Qiao, B., Li, L., Zhang, X., Lin, Q., & Lyu, M. (2020). AIOps innovations of incident management for cloud services. In Proceedings of the Cloud Intelligence Workshop at the 34th AAAI Conference on Artificial Intelligence.
2. Chen, Z., Kang, Y., Li, L., Zhang, X., Zhang, H., Xu, H., Zhou, Y., Yang, L., Sun, J., Xu, Z., Dang, Y., Gao, F., Zhao, P., Qiao, B., Lin, Q., Zhang, D., & Lyu, M. (2020). Towards intelligent incident management: Why we need it and how we make it. In Proceedings of the 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (pp. 1494–1505).



3. Sasi Kumar Kolla. (2021). Adaptive Graph-Based Framework for Secure and Efficient Distributed Data Processing. *Journal of Computational Analysis and Applications (JoCAAA)*, 29(6), 1519–1523. Retrieved from <https://www.eudoxuspress.com/index.php/pub/article/view/5860>
4. Dayal, R., Vijayakumar, V., Kushwaha, R. C., Kumar, A., Kumar, V. D. A., Kumar, A., Varadarajan, V., Kommers, P., Piuri, V., & Subramaniaswamy, V. (2020). A cognitive model for adopting ITIL framework to improve IT services in Indian IT industries. *Journal of Intelligent & Fuzzy Systems*, 39(6).
5. Serrano, J., Faustino, J., Adriano, D., Pereira, R., & Mira da Silva, M. (2021). An IT service management literature review: Challenges, benefits, opportunities and implementation practices. *Information*, 12(3), Article 111.
6. RAMI, R. M. K. (2020). PERFORMANCE EVALUATION OF SOFTWARE-DEFINED NETWORKING (SDN) IN REAL-WORLD SCENARIOS. *INTERNATIONAL JOURNAL*, 9(3), 1665-1671.
7. Kolla, S. H. (2021). Rule-based automation for IT service management workflows. *Online Journal of Engineering Sciences*, 1(1), 1–14.
8. Spring, M., Faulconbridge, J., & Sarwar, A. (2022). How information technology automates and augments processes: Insights from artificial-intelligence-based systems in professional service operations. *Journal of Operations Management*, 68(6–7), 592–618.
9. Satriya Wibawa, N. M., & Ramantoko, G. (2022). Business process analysis of cloud incident management service with activity assignment: A case of PT. XYZ. *Journal of Business Management and Accounting*, 12(1), 51–80.
10. Mashetty, S., & Mukesh, A. (2022). Application of Machine Learning in Loan Default Forecasting. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 3(4), 123-132. <https://doi.org/10.63282/3050-9262.IJAIDSML-V3I4P114>
11. Machine learning-based zero-touch network and service management: A survey. (2022). *Digital Communications and Networks*, 8(2), 105–123.
12. Research framework for determining how artificial intelligence enables information technology service management for business model resilience. (2021). *Sustainability*, 13(20), Article 11496.
13. Sreerangapuri, A. (2021). IT service management transformation: Implementing AIOps for proactive service delivery. *NeuroQuantology*, 19(10).



14. Sreenivasulu Naidu Yalakaturi. (2022). Rule-Based Data Quality Monitoring for Multi-System Healthcare Information. *International Journal of Medical Toxicology and Legal Medicine*, 25(3 and 4), 23–31. Retrieved from <https://ijmtlm.org/index.php/journal/article/view/1504>
15. Chen, Z., et al. (2020). Research on intelligent incident management in cloud services. *Proceedings of the ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*.
16. Dayal, R., et al. (2020). Cognitive approaches to IT service management and ITIL adoption. *Journal of Intelligent & Fuzzy Systems*.
17. Serrano, J., et al. (2021). IT service management implementation practices and organizational challenges. *Information*, 12(3).
18. Manukonda, K. R. R. (2021). ENHANCING SOFTWARE QUALITY ASSURANCE THROUGH MACHINE LEARNING: TECHNIQUES, APPLICATIONS, AND CASE STUDIES. *JOURNAL OF BASIC SCIENCE AND ENGINEERING*, 18(1), 202-210.
19. Spring, M., et al. (2022). Artificial intelligence in professional service operations: Automation and augmentation. *Journal of Operations Management*, 68(6–7), 592–618.
20. Kolla, S. H. (2021). Rule-based automation for IT service management workflows. *Online Journal of Engineering Sciences*, 1(1), 1–14.
21. Satrya Wibawa, N. M., & Ramantoko, G. (2022). Business process analysis of cloud incident management service with activity assignment. *Journal of Business Management and Accounting*, 12(1), 51–80.
22. Chen, Z., et al. (2020). AIOps innovations in cloud incident management and service reliability. *Cloud Intelligence Workshop at the AAAI Conference on Artificial Intelligence*.
23. Research framework for determining how artificial intelligence enables information technology service management for business model resilience. (2021). *Sustainability*, 13(20), Article 11496.
24. Machine learning-based zero-touch network and service management: A survey. (2022). *Digital Communications and Networks*, 8(2), 105–123.
25. Rule-based automation for IT service management workflows. (2021). *Online Journal of Engineering Sciences*, 1(1), 1–14.
26. AIOps innovations of incident management for cloud services. (2020). *Proceedings of the Cloud Intelligence Workshop at the AAAI Conference on Artificial Intelligence*.



27. Towards intelligent incident management: Why we need it and how we make it. (2020). Proceedings of the ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering.
28. Anil Lokesh Gadi, “Connected Mobility Platforms and Their Impact on Urban Transport Planning and Service Design,” International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering (IJIREEICE), DOI: 10.17148/IJIREEICE.2019.71207
29. A cognitive model for adopting ITIL framework to improve IT services in Indian IT industries. (2020). Journal of Intelligent & Fuzzy Systems, 39(6).
30. How information technology automates and augments processes: Insights from artificial-intelligence-based systems in professional service operations. (2022). Journal of Operations Management, 68(6–7), 592–618.
31. Business process analysis of cloud incident management service with activity assignment: A case of PT. XYZ. (2022). Journal of Business Management and Accounting, 12(1), 51–80.