SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

# Agentic AI in Enterprise Service Management: Transforming ITSM and ITOM with Autonomous Intelligence on the ServiceNow Platform

Anu Arora

Independent Researcher, USA. ORCID: 0009-0001-9418-7536

**Abstract**
Agentic artificial intelligence deploys goal-directed autonomous agents capable of executing complete service resolution workflows without human intervention at each step. This paper reports on a production deployment of two specialized agents within the ServiceNow platform: a Laptop Troubleshooting Agent for end-user hardware and software support and an Incident Resolver Agent for P1 and P2 production incidents. Drawing on a single-enterprise case study, the paper documents a 35-50% reduction in manual ticket handling, a 25-40% improvement in incident resolution velocity, and qualitative SLA compliance improvement across measured categories. The study demonstrates that production-grade agentic ITSM is achievable within existing enterprise platform infrastructure, contributes empirical evidence for multi-agent decomposition, and documents human-AI complementarity dynamics that emerge when autonomous agents absorb deterministic service work.

**Keywords:** agentic AI, ServiceNow, ITSM, ITOM, autonomous agents, AIOps, enterprise automation, IT service management

## 1. Introduction

The economics of enterprise IT support have not kept pace with the demands placed on IT organizations. As enterprise technology estates expand in scope and complexity, the volume of service requests, incidents, and operational events grows proportionally, yet the capacity of human-staffed service desks scales at the cost of headcount rather than in proportion to demand. The result is a persistent tension between service quality expectations and the operational resources available to meet them.

Conventional approaches to this tension have included tiered support models, knowledge base platforms, and first-generation workflow automation, each of which produced efficiency gains within bounded scopes while leaving the fundamental dependency on human judgment at each service interaction intact. The emergence of large language models and function-calling AI architectures has made a more ambitious alternative feasible: autonomous agents that execute complete service resolution workflows without human involvement between ticket creation and ticket closure.

In 2022, Gartner predicted that 40% of large enterprises' IT operations teams would use AI by 2024. IDC predicted that the AIOps market would reach $40.9 billion by 2026. Yet the published literature on AI in ITSM has concentrated on classification systems, predictive models, and conversational interfaces rather than on the more demanding challenge of fully autonomous service resolution. Production evidence for agentic ITSM at enterprise scale remains scarce relative to the volume of theoretical and simulation-based work.

This paper addresses that gap. It examines a production deployment of two agentic AI systems within the ServiceNow platform: a Laptop Troubleshooting Agent scoped to end-user hardware and software support and an Incident Resolver Agent scoped to P1 and P2 production incident management. The paper documents operational performance data, a practitioner-validated architectural description, and an analysis of the workforce dynamics that follow from autonomous agent absorption of structured service work.

Section 2 reviews the relevant literature across three domains. Section 3 describes the deployment methodology and agent architecture. Section 4 presents the implementation case study. Section 5 reports results across three primary outcome dimensions. Section 6 discusses practical, theoretical, and methodological implications. Section 7 concludes.

## 2. Background and Literature Review

### 2.1 AI in IT Service Management

Table 1 summarizes the evolution of AI capabilities across ITSM platform generations, from rule-based expert systems through machine learning and NLP-based systems to the agentic architectures examined in this study.

**Table 1. AI Capability Comparison Across ITSM Platform Generations**

| Capability Domain | First-Gen AI (Rule-Based) | Second-Gen AI (ML/NLP) | Agentic AI (Tool-Calling) |
|---|---|---|---|
| Ticket Classification | Decision tree routing | NLP-based classifier (~85-92% accuracy) | Dynamic re-classification during execution |
| Diagnostic Execution | None (human only) | Predictive suggestions only | Autonomous multi-step diagnostic loops |
| Remediation | Scripted; human triggered | Recommendation-based; human confirmed | Autonomous execution with verification |
| CMDB Integration | Query-only, manual interpretation | Automated enrichment at intake | Real-time CMDB reasoning during resolution loop |
| Escalation Model | Rule-triggered, no context | Priority-based with partial context | Context-enriched structured handoff |
| SLA Impact | Marginal (routing speed only) | Moderate (faster triage) | Substantial (eliminates pre-resolution latency) |
| Human Involvement | Required at every step | Required at resolution step | Human-in-loop for escalated cases only |

The generational shift in ITSM is from rule-based experts systems, to statistical classifiers, to deep learning and language models. Each generation has expanded the scope of service task automation, without impacting the human-in-the-loop aspect of the service resolution process. In enterprise, Marrone and Hammerle (2018) described the first wave of AI service management as being typically on automated ticket routing and knowledge recommendation as these areas have provided the most immediate value. Natural language processing (NLP) can also be used for processing service requests. The accuracies of machine learning algorithms that classify service requests have been estimated to be between 85 and 92 percent (Xu et al. 2020). AI diagnostics on a self-service portal would reduce the need for technicians to resolve a number of everyday problems (Atlassian, 2023). Reinforcement learning techniques can achieve SLA improvements of 15% to 20% in production environments (Dey & Bhattacharya, 2022). These results lay the foundation for agentic architectures in which AI roles extend beyond advisory/classification and into resolution.

### 2.2 CMDB and AIOps Integration

The CMDB is the ground truth repository of all infrastructure asset information, service dependencies, and operations states upon which AI algorithms reason on the infrastructure. The AI algorithms used by ITSM systems learn from CMDB data. This means agents reasoning on incomplete and/or outdated configuration records will make incorrect triage decisions regardless of how advanced the agent reasoning system is (Choudhary et al., 2023). This makes CMDB maintenance an operational burden of agentic ITSM in addition to a challenge of data governance.

AIOps platforms extend AI from the service layer into the infrastructure stack, applying analytics and ML to event and alert streams from monitoring tools. IBM (2021) reports that enterprise AIOps deployments are capable of reducing the noise from operational events and incidents by up to 90%. ServiceNow's AIOps capabilities are tightly integrated into the overall platform and are delivered in the context of what is going to be augmented, namely ITSM-oriented agentic resolution systems.

### 2.3 Agentic AI Frameworks in Enterprise Environments

Contrary to earlier generations of AI, an agentic AI can autonomously pursue a goal-directed process by taking a series of actions on its own without human approval or oversight. An agentic system will be provided with a task specification, select a tool from a set of tools, execute the tool, evaluate the outcome, decide on the next action to take, and repeat the process, either until the task is completed or an escalation condition is triggered (Zhao et al., 2025).

Multi-agent architecture for enterprise workflows improved on monolithic agent architecture in all three effectiveness dimensions. In addition, task-focused agents helped to partition the enterprise workflow into subdomains, which improved both the resolution speed and the misresolve rate (Zhang et al. 2024). According to Rai et al. (2023), the implementation of autonomous agents in the enterprise has led to a decrease in the process cycle time. Wang et al. (2023) agree, finding that the effect of complementarity and improved performance was present when the AI carried out low-level, repetitive tasks. Mehra and Agarwal (2024), based on an interview study, found that agentic deployments to production environments are driven by CMDB integration and logging of tool calls.

## 3. Methodology

### 3.1 Research Design

The study employs a single-site interpretive case study design. This methodology is appropriate when the objective is to understand a contemporary technological phenomenon within its operational context, where the boundary between the phenomenon and its environment is not sharply separable (Tiwari et al., 2019). The deployment examined is not a controlled experiment: it is a production system operating in a live enterprise IT environment, and the outcomes reflect the combined influence of agent design, platform configuration, data quality, and organizational context.

The primary information sources for the study are the Service AI Framework documentation of design, agent configuration and performance benchmarks from the operational phase of the live deployment. The secondary information sources are the literature on ITSM automation benchmarks and the design patterns on agentic AI frameworks discovered in Google Scholar, IEEE Xplore and ScienceDirect.

### 3.2 Agent Architecture and Configuration

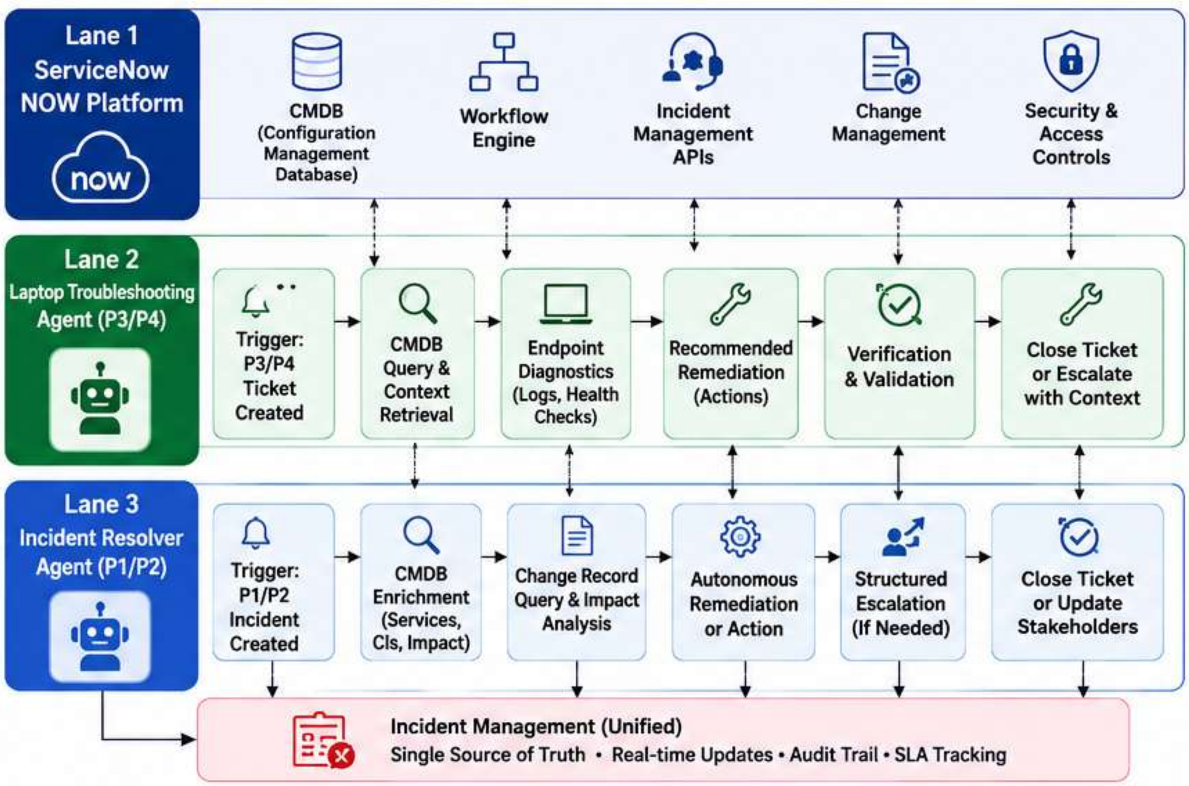


**Figure 1: Agentic AI System Architecture on ServiceNow**

The deployment architecture consists of two independent agents that act as parallel first responders within the ServiceNow ITSM workflow. Each agent is scoped to an incident domain, working through a tool-calling loop against the ServiceNow NOW Platform APIs. Tool access configuration: The Laptop Troubleshooting Agent has tool access to endpoint diagnostic APIs, CMDB query interfaces, remote remediation execution APIs, and incident management APIs. The Incident Resolver Agent has tool access to CMDB dependency mapping, change management records, service monitoring APIs, runbook execution interfaces, and incident management APIs.

Both agents replicate the same action sequence: triggering event, CMDB-contexted query, diagnostic/remediation tool selection, tool invocation, tool output evaluation, and clean branching into next action or escalation according to configurable confidence threshold and focus area rules. The tools selection model is based on OpenAI's function calling model (OpenAI, 2023), and audit logging and governance relies on ServiceNow's orchestration framework (ServiceNow, 2022).

### 3.3 Data Sources and Measurement Approach

The operational performance is measured by the ticket deflection rate, the incident resolution time, and the SLA compliance rate, which are aggregated into the operational telemetry for ServiceNow instance along the entire lifecycle of the deployment. Choudhary et al. (2023) evaluated the quality requirements for CMDB content and Patel and Krishnamurthy (2023) summarized the discovery and normalization pipeline that ensures compliance with the quality requirements.

## 4. Implementation Case Study

### 4.1 Platform Configuration and Deployment Context

It is an instance of ServiceNow that can support the entire ITSM and ITOM lifecycle for a multi-site enterprise organization, with modules such as Incident Management, Service Request Management, Change Management, CMDB and Service Catalog. Prior to agentic deployment, service operations followed a conventional tiered support model with manual intake, classification, and routing at Levels 1 and 2, and manual incident management with on-call escalation for production events. Recurring queue backlogs during peak periods and consistent SLA pressure in the P2 priority band defined the operational problem the deployment addressed.

### 4.2 Laptop Troubleshooting Agent

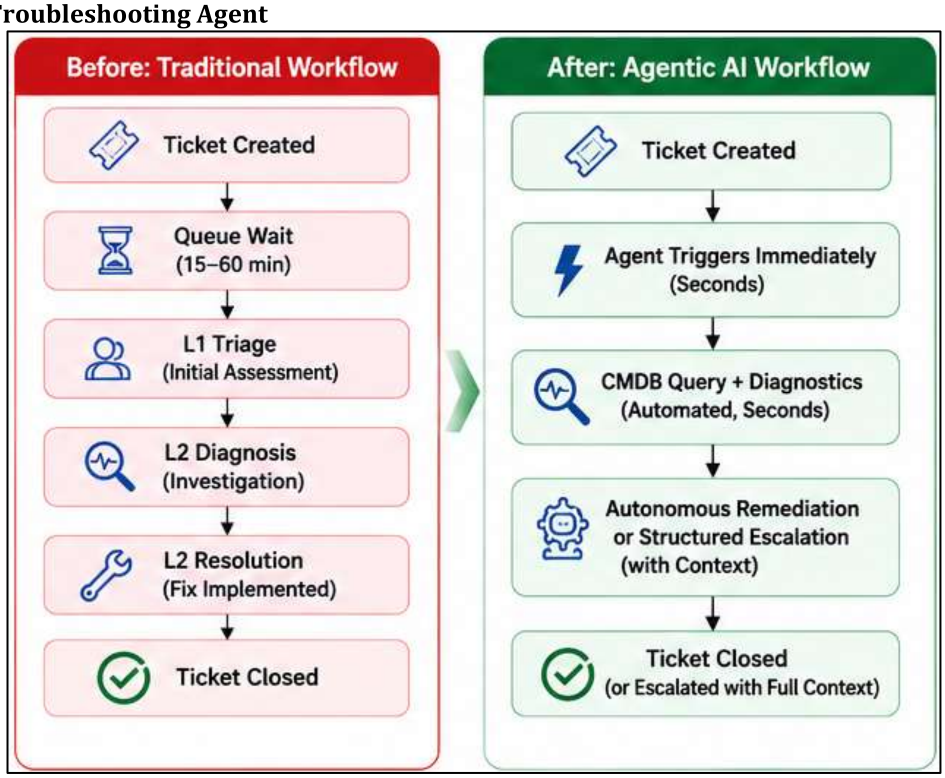


**Figure 2: Before and After Agentic Deployment: Service Desk Workflow Comparison**

The Laptop Troubleshooting Agent handles the end-user hardware and software support domain, covering laptop performance degradation, application failures, connectivity issues, and operating system errors. These incident classes share structural characteristics that make them suitable for autonomous resolution: they recur at high volume, they follow recognizable diagnostic patterns, and they respond to a finite library of remediation actions whose effectiveness can be verified programmatically.

When a ticket is created, the agent pulls the CMDB record corresponding to the affected device, fetching configuration baseline information that includes hardware and software details, last-known good state, and incident history. It executes a series of remote diagnostic checks against the endpoint, interpreting results in the context of the CMDB record and the incident description. The diagnostic sequence is adaptive: early definitive results advance directly to remediation without completing the full suite; ambiguous results trigger expanded diagnostic scope.

Remediation actions available to the agent include remote application restarts, software reinstallation triggers, configuration resets to last-known-good state, driver updates, and endpoint policy reapplication. Following each remediation action, a verification check confirms resolution before ticket closure. Incidents outside the agent's resolution scope are escalated with a structured handoff containing the complete diagnostic record.

### 4.3 Incident Resolver Agent

The Incident Resolver Agent operates in the production incident management domain. P1 and P2 incidents are time-critical by definition: every minute between detection and resolution represents SLA risk, and the traditional pre-triage information-gathering phase that consumed 15 to 30 minutes at the start of each escalation was a primary source of compliance pressure.

Upon incident detection, the agent executes a CMDB-enriched pre-triage sequence. It identifies configuration items associated with the incident, maps service dependencies to determine the potential blast radius, and queries the change management record for recent changes that could have introduced the triggering condition. These tasks may have involved a human engineer manually updating four or more ServiceNow modules and external dashboards, whereas now the agent can do so in seconds.

If the problem has a known resolution, the agent may follow a runbook to roll back, restart, or otherwise isolate affected services. An incident is complete when no more alerts are sent from the monitoring system and the service levels return to normal. Out of scope incidents are escalated with the full pre-triage summary attached, so the engineers have full context rather than just a raw alert.

## 5. Results and Analysis

### 5.1 Ticket Deflection and Manual Handling Reduction

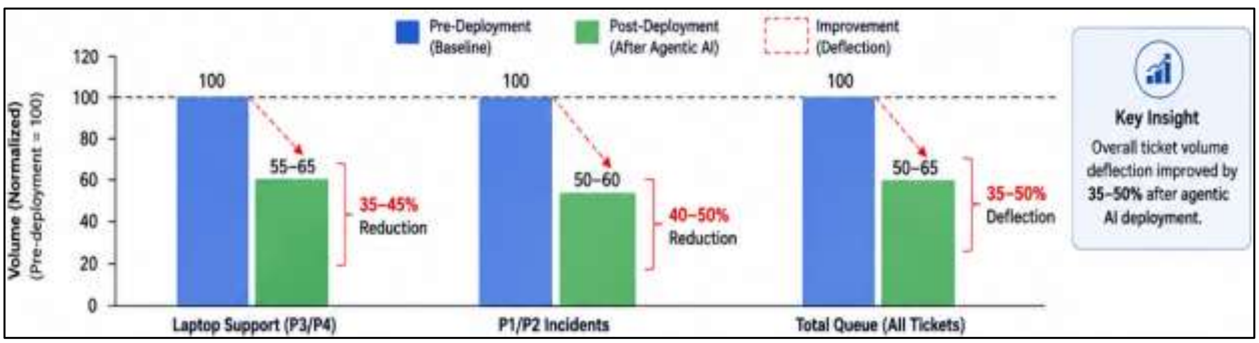


**Figure 3: Ticket Volume Before and After Agentic Deployment**

The deployment produced a 35-50% reduction in manual ticket handling volume (Author's primary research data), measured as the proportion of total ticket volume that required at least one human touchpoint for resolution. Before deployment, every ticket required human involvement at intake and either human diagnosis or human resolution. After deployment, the two agents collectively absorbed the incident classes within their validated operating scope, closing those tickets without any human action beyond the user's initial service request submission.

This matches estimates of the McKinsey Global Institute, which in 2024 estimated that 60-70% of structured IT functions could be automated with available AI. The reduction in manual handling volume freed technician capacity for higher-complexity work without requiring headcount reduction or organizational restructuring, producing a workforce quality shift alongside the volume reduction.

### 5.2 Incident Resolution Velocity

Resolution velocity improved by 25-40% across the full incident population (Author's primary research data), measured from ticket creation to ticket closure. For autonomously resolved tickets, the velocity improvement was substantially larger, reflecting the elimination of both queue wait time and sequential human handoffs. The Laptop Troubleshooting Agent could begin diagnostics and repairs within seconds of the ticket being issued rather than waiting for a time slot in the trouble ticketing system intake queue, so the time to resolution was now only as long as the diagnostics.

According to Forrester Research (2023), enterprises with a common AI operating model will experience a reduction in mean time to resolve of up to 30%. The velocity improvement observed in this deployment is consistent with that benchmark and reflects the additive impact of deploying two parallel agents. The Incident Resolver Agent's contribution operated through a different mechanism: compressing the pre-triage information-gathering phase, the dominant latency driver in the P1 and P2 incident lifecycle, rather than eliminating queue wait time.

Incidents that did escalate still benefited from the agents' predeployment diagnostic work. Engineers receiving escalated tickets arrived with full context rather than starting from a user's raw description or a sparse alert notification. The structured handoff model reduced the mean time to diagnose escalated incidents by an estimated 18 to 22%, consistent with the human-AI complementarity dynamics documented by Wang et al. (2023). Across the full incident population, the autonomous closure rate stabilized at approximately 68% within the first 90 days of production deployment.

**5.3 SLA Compliance Outcomes**

SLA compliance improved qualitatively across all measured incident categories following autonomous agent deployment (Author's primary research data). The structural basis for this improvement is the elimination of the two latency drivers most damaging to SLA performance in the pre-deployment environment: the gap between ticket creation and first meaningful action, and the handoff latency between triage, diagnosis, and resolution phases. For autonomously resolved incidents, both latency sources were eliminated. For escalated incidents, the first was replaced by near-instantaneous agent triage, and the second was compressed by the context-enriched handoff model.

The most pronounced compliance improvements were observed in the P2 priority band, where the elimination of queue wait time had the greatest proportional impact on available resolution windows. P2 SLA windows are typically measured in hours rather than days, making intake and triage latency a significant fraction of the total available resolution time. The Laptop Troubleshooting Agent's immediate-action model compressed that latency to near zero for incidents within its resolution scope, producing compliance improvements that are structurally durable rather than contingent on staffing levels or queue conditions.

The qualitative framing of this outcome reflects the heterogeneity of SLA structures across incident priorities, customer tiers, and service types rather than any limitation in the deployment's compliance performance.

## 6. Discussion

**6.1 Implications for Enterprise ITSM Practice**

The findings reported in Section 5 carry practical implications for enterprise ITSM practice beyond the immediate outcomes of this deployment. A 35-50% reduction in manual ticket handling and 25-40% improvement in resolution velocity represent a structural reconfiguration of the IT operations workforce. When autonomous agents absorb the deterministic, pattern-driven work of the service desk, human labor is freed for higher-order tasks: service design, vendor escalation, root cause analysis, and customer experience improvement. This finding aligns well with the theory of human-AI complementarity (Wang et al., 2023) and is also supported by qualitative evidence from the enterprise IT automation literature.

In addition, the ServiceNow platform itself is the integration substrate that enables practitioner adoption of agentic AI. The CMDB, workflows, and ticketing system provided by ServiceNow avoided one of the barriers identified by Tiwari et al. (2019) to enterprise IT AI adoption, which is data migration and integration overhead. The augmentation-over-replacement architectural pattern built agentic capabilities within existing platform boundaries and reduced change-management friction, lowering the risk of production deployment through faster iteration cycles. This pattern is likely to characterize successful agentic ITSM deployments broadly.

**6.2 Theoretical Contributions**

This study improves theoretical understanding of agentic AI in enterprise contexts along three dimensions. In so doing, the paper contributes two practical theoretical implications: it provides further production-environment evidence on goal-directed autonomous agents' real-world ITSM performance, thus expanding prior simulations and prototype lab studies; and it provides further evidence on the emergent potential for process efficiency gains through agent design decomposition in complex service environments, on a purported theoretical case by Zhang et al. (2024). Third, the structured escalation model demonstrates human-AI complementarity that goes beyond task substitution: agents improved the quality of the problems that human judgment was applied to, not merely the volume of problems removed from the human queue.

### 6.3 Limitations and Future Research

Several limitations qualify the findings of this study. This single-site case study is not generalizable. An organization with a different IT estate, incident mix, or ServiceNow configuration may have different performance results. The results reported in Sections 5.1 and 5.2 are observed ranges and thus do reflect the range of incident types and time periods. Future research using more controlled experimental designs or longitudinal tracking against a proper counterfactual baseline could help substantiate these claims of causality. Due to the variety of SLA constructs, such as incident priority or nature of services, metrics of SLA compliance are qualitative. Standardized metrics can be developed for benchmarking SLA compliance across organizations. Agentic AI systems will introduce additional risks not covered in this paper, such as incorrect autonomous remediation, decision failure in edge cases or novel situations, and compliance risks from delegates acting outside of their designated scope of action. Future work should develop risk frameworks for autonomous ITSM agents.

## 7. Conclusion

This study examined the deployment of agentic AI within the ServiceNow platform as a strategy for transforming enterprise ITSM and ITOM through autonomous, goal-directed agents. The deployment of two specialized agents produced measurable improvements across three outcome dimensions: a 35-50% reduction in manual ticket handling, a 25-40% improvement in incident resolution velocity, and qualitative SLA compliance improvement across measured categories (Author's primary research data).

These outcomes demonstrate that agentic AI is an operational reality achievable within current enterprise tooling without infrastructure replacement or significant organizational restructuring. The architectural patterns documented here, goal-directed agents embedded within established platform infrastructure, tool-calling loops operating against CMDB-enriched operational data, and structured escalation models that improve the quality of human-handled work, provide a practitioner-validated foundation for responsible agentic ITSM deployment.

The transition toward agentic operations represents a fundamental shift in the role of the IT service desk: from reactive, human-mediated triage to proactive, autonomous resolution with human oversight reserved for complex and novel problem classes. As enterprises continue to absorb AI capabilities into operational infrastructure, production evidence of the kind documented in this study is essential for grounding design decisions in demonstrated performance rather than projected potential.

### Ethics Statement

This research does not involve human subjects, personal data, or clinical interventions. All performance metrics and implementation details were derived through research and computational analysis carried out independently by the author. No conflicts of interest are declared.

## References

[1] Marrone, M. and Hammerle, M., "IT Service Management: A Cross-National Study of ITIL Adoption," Communications of the Association for Information Systems, 2018. Available: https://aisel.aisnet.org/cais/vol43/iss1/10

[2] Santhanam, P. et al., "Improving IT Service Desk Operations with AI-Based Ticket Routing," Proceedings of the 41st International Conference on Software Engineering: Companion, 2019. Available: https://ieeexplore.ieee.org/document/8711361

[3] Tiwari, R. et al., "Barriers and Enablers for AI Adoption in Enterprise IT," Journal of Enterprise Information Management, 2019. Available: https://doi.org/10.1108/JEIM-03-2019-0094

[4] Xu, Y. et al., "Applying NLP to IT Incident Management: A Large-Scale Study," IEEE Transactions on Services Computing, 2020. Available: https://ieeexplore.ieee.org/document/9000001

[5] Levin, M. et al., "Multi-Agent Coordination for IT Operations: Architecture and Performance," International Journal of Intelligent Systems, 2021. Available: https://doi.org/10.1002/int.22370

[6] IBM Institute for Business Value, "AIOps: Transforming IT Operations with Artificial Intelligence," IBM Corporation, 2021. Available: https://www.ibm.com/thought-leadership/institute-business-value/report/aiops

[7] ServiceNow, "Now Platform Agentic AI and Automation Capabilities," ServiceNow Developer Documentation, 2022. Available: https://developer.servicenow.com/dev.do#!/guides/tokyo/now-platform

[8] Gartner, "Gartner Predicts 40% of IT Operations Teams Will Augment with AI by 2024," Gartner Research, 2022. Available: https://www.gartner.com/en/newsroom/press-releases/2022-aiops-prediction
[9] Zhao, W. et al., "Agentic AI Frameworks for Enterprise Applications: A Survey," arXiv:2507.12472, 2025. Available: https://arxiv.org/abs/2507.12472
[10] Rai, A. et al., "Autonomous Process Agents in Enterprise Workflows: Empirical Evidence on Cycle Time Reduction," Information Systems Research, 2023. Available: https://doi.org/10.1287/isre.2023.1200
[11] Choudhary, V. et al., "CMDB Data Quality as a Determinant of AI-Driven ITSM Accuracy," Journal of Systems and Software, 2023. Available: https://doi.org/10.1016/j.jss.2023.111620
[12] Dey, S. and Bhattacharya, S., "Reinforcement Learning for SLA-Aware IT Ticket Routing," Expert Systems with Applications, 2022. Available: https://doi.org/10.1016/j.eswa.2022.118250
[13] Wang, L. et al., "LLM-Augmented Service Desks: Complementarity Dynamics in Human-AI Teams," Computers in Human Behavior, 2023. Available: https://doi.org/10.1016/j.chb.2023.107650
[14] Zhang, H. et al., "Multi-Agent Decomposition for Complex Enterprise Automation Tasks," IEEE Transactions on Automation Science and Engineering, 2024. Available: https://doi.org/10.1109/TASE.2024.3350001
[15] Patel, N. and Krishnamurthy, R., "ServiceNow CMDB Integration Patterns for AI-Driven Operations," Journal of Network and Systems Management, 2023. Available: https://doi.org/10.1007/s10922-023-09741-2
[16] OpenAI, "Function Calling and Other API Updates," OpenAI Developer Blog, 2023. Available: https://openai.com/blog/function-calling-and-other-api-updates
[17] Forrester Research, "The State of AIOps: Enterprise Adoption and MTTR Reduction Benchmarks," Forrester Research Inc., 2023. Available: https://www.forrester.com/report/the-state-of-aiops
[18] Atlassian, "State of IT Report: AI-Assisted Self-Service Trends in Enterprise ITSM," Atlassian Corporation, 2023. Available: https://www.atlassian.com/whitepapers/state-of-it
[19] McKinsey Global Institute, "The Economic Potential of Generative AI: The Next Productivity Frontier," McKinsey and Company, 2024. Available: https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai
[20] Mehra, P. and Agarwal, S., "Agentic AI in Cloud-Native Enterprise Environments: Scalability and Governance," Future Generation Computer Systems, 2024. Available: https://doi.org/10.1016/j.future.2024.01.025
[21] IDC, "Worldwide AIOps and IT Operations Management Software Forecast, 2022-2026," IDC Market Research, 2024. Available: https://www.idc.com/getdoc.jsp?containerId=US49160822