International Journal of Science, Research and Technology (IJSRAT)

|Bimonthly | Peer-Reviewed | Scholarly Journal|

|ISSN 2454-3853|www.ijsrat.com|editor@ijsrat.com|Vol.9, Issue 1, January-February 2026|

DOI: 10.15662/IJSRAT.2026.0901007

Operational Intelligence for SAP: How AI Agents Transform Incident Response and System Health

Anuradha Karnam

Principal Cloud Solution Architect, Microsoft Corporation, USA

Publication History: 11-12-2025 (Received); 30-12-2025 (Revised); 9-1-2026 (Accepted); 18-1-2026 (Published).

ABSTRACT: For over two decades, the management of deeply integrated enterprise ecosystems has relied on the rigid, human-in-the-loop governance of traditional IT Service Management (ITSM) frameworks a paradigm that sprawling SAP and hybrid cloud infrastructures have rendered functionally obsolete. The industry's subsequent pivot to legacy AIOps attempted to stem this operational bleeding through deterministic log parsing and alert suppression, which has become the approach for a decade. Such deterministic tools operate as narrow, procedural toys that flag anomalies but fundamentally lack the distributed cognitive architectures required for automated remediation. But what, then, is the actual operational utility of merely observing a failure without the algorithmic capacity to heal it? To move beyond the methodological navel-gazing that currently plagues the field, this research proposes a foundational reorientation: an Integrated Incident Response Model (IIRM) wherein goal-driven Agentic AI autonomously ingests unstructured SAP application logs and AWS telemetry, utilizing contextual multi-armed bandit optimization to dynamically negotiate root cause analysis. Empirical evaluation within a rigorously simulated, high-volatility hybrid cloud deployment demonstrates that these communicating agents profoundly reduce Mean Time to Resolution (MTTR) and maintain high diagnostic accuracy, entirely decoupling system recovery from human latency. While mathematically provable guardrails remain an absolute necessity to prevent catastrophic compliance failures, this transition from reactive monitoring to proactive algorithmic agency capable of reflecting, learning, and improving over time finally provides the architecture required to redefine the twenty-year arc of enterprise system resilience.

KEYWORDS: Operational Intelligence, Sap Systems, Incident Response, Agentic AI, Autonomous IT Operations, Automated Remediation, Mean Time To Resolution (MTTR)

I. INTRODUCTION

For the better part of two decades, the management of deeply integrated enterprise ecosystems has relied on rigid, human-in-the-loop oversight. We built entire disciplines around operational intelligence and IT Service Management (ITSM), codifying best practices into monolithic frameworks like ITIL v3 to establish a predictable, standardized governance. This approach served its historical purpose: it attempted to transition IT from a chaotic, reactive firefighting team into a reliable service provider aligned with business needs.

However, as modern SAP landscapes sprawl across hybrid cloud infrastructures integrating diverse modules from financial accounting and human capital management to supply chain logistics they generate a velocity of telemetry that renders manual incident response not merely inefficient, but functionally obsolete. The industry's recent pivot to AIOps (AI for IT Operations) was intended to stem this bleeding by automating ticket classification and parsing logs.

The Limitations of Deterministic AIOps in Automated Remediation

Current AIOps implementations remain fundamentally deterministic components with a severely limited scope. They are exceptionally adept at parsing unstructured data and flagging anomalies to reduce alert fatigue, but they lack the distributed cognitive architectures required for autonomous IT operations. They alert us to the fire, but they cannot hold the hose. We must ask: what is the actual utility of an intelligent monitor if the remediation still requires a Level-3 support engineer to manually execute the fix? These legacy machine learning tools operate as narrow, procedural toys rather than robust operational frameworks [17]. The persistent gap lies in bridging system observability with automated remediation a cognitive leap that requires moving from static, pattern-matching algorithms to goal-driven Agentic AI capable of dynamic root cause analysis, inter-agent communication, and true predictive maintenance [3, 15].

International Journal of Science, Research and Technology (IJSRAT)

|Bimonthly | Peer-Reviewed | Scholarly Journal|

|ISSN 2454-3853|www.ijsrat.com|editor@ijsrat.com|Vol.9, Issue 1, January-February 2026|

DOI: 10.15662/IJSRAT.2026.0901007

The Shift Toward Distributed Agentic AI and the IIRM

To move beyond the methodological navel-gazing that currently plagues operational intelligence, we propose a fundamental reorientation of system health management. The integration of Agentic AI systems characterized by goal decomposition, distributed intelligence, and contextual adaptation provides the necessary architecture for a robust Integrated Incident Response Model (IIRM) [4, 5]. Unlike basic automation scripts, these distributed agents continuously ingest raw SAP application logs and AWS CloudWatch telemetry, delegating diagnostic tasks across specialized nodes when anomalous thresholds are breached [1]. By modeling incident resolution as a contextual multi-armed bandit optimization, the framework balances the exploitation of known fixes with the exploration of novel remediation paths. Preliminary evaluations of this architecture demonstrate profound reductions in Mean Time to Resolution (MTTR) and significant improvements in diagnostic accuracy, allowing systems not only to act but to reflect, learn, and improve over time without continuous human supervision [9].

Addressing Data Integrity and Regulatory Compliance

Yet, we must be clear-eyed about the boundaries of this approach. Agentic AI is not a panacea, nor does it operate in a vacuum. True AI Operationalization spans an entire lifecycle from the creation and training of algorithms to their deployment, monitoring, curation, and eventual retirement. This framework assumes a high degree of foundational data integrity; deploying autonomous agents into poorly configured SAP environments that generate noisy, unstandardized logs, or where synchronization disparities exist between SAP and non-SAP systems, will inevitably cause the system to hallucinate relationships that do not exist. Furthermore, automated remediation within critical financial or human resource modules carries inherent regulatory and compliance risks that algorithmic efficiency cannot simply wave away. Executing autonomous database modifications in a highly regulated enterprise environment requires mathematically provable guardrails [18]. If the system cannot guarantee compliance, its speed is irrelevant.

Ultimately, the transition from reactive monitoring to proactive, agentic intervention profoundly redefines enterprise resilience. To understand how this proposed IIRM departs from historical precedents, we must first examine the cyclical evolution of incident response frameworks and the critical distinction between basic automation and true algorithmic agency.

II. LITERATURE REVIEW

If we examine the trajectory of operational excellence over the past two decades, a distinct, cyclical pattern emerges. Early efforts focused entirely on process management via ITIL and ITSM frameworks [20]. We codified the five core lifecycle phases Service Strategy, Service Design, Service Transition, Service Operation, and Continual Service Improvement establishing a rigid "Way of Working" that governed human behavior but entirely lacked computational foresight. We built governance, not intelligence. Over the last decade, the literature reflects a desperate pivot to AIOps, integrating deterministic machine learning to parse unstructured data and reduce alert fatigue [13, 14]. This became the approach for managing sprawling cloud infrastructures.

Given this dead end, we must ask: what is the actual contribution of these legacy AIOps tools? The prevailing research champions them as revolutionary, pointing to marginal improvements in anomaly detection and automated ticket classification. Yet, they remain largely glorified pattern matchers. They operate as narrow, procedural toys that wait for human intervention to execute automated troubleshooting. The fundamental flaw in the current literature is the persistent conflation of automation with autonomy [8].

Evaluating the Gap Between Basic Automation and True Autonomy

Recognizing this conflation, recent scholarship has finally begun to draw a critical and long overdue distinction between basic AI agents and true Agentic AI [7]. The former are deterministic components constrained by single-task scripts; they fail spectacularly when confronted with the multi-stage, cascading anomalies typical of hybrid SAP and AWS environments [10, 11]. Conversely, Agentic AI represents distributed intelligence, characterized by goal decomposition, inter-agent communication, and contextual adaptation. While analyses indicate that a distinct connection exists between artificial intelligence and operational excellence specifically in leveraging automated intelligent algorithms to find patterns and automatically select optimal operational paths the broader literature remains heavily bogged down by theoretical orchestration rather than empirical deployment. We are drowning in abstract frameworks while enterprise systems continue to rely on manual, human-driven root cause analysis.

This theoretical stagnation is particularly glaring in the context of SAP ecosystems. While general cloud infrastructure research frequently extols the virtues of intelligent automation, integrating these autonomous frameworks within highly regulated, tightly coupled enterprise resource planning systems remains an unresolved tension. Automated remediation

International Journal of Science, Research and Technology (IJSRAT)
|Bimonthly | Peer-Reviewed | Scholarly Journal|
|ISSN 2454-3853|www.ijsrat.com|editor@ijsrat.com|Vol.9, Issue 1, January-February 2026|
DOI: 10.15662/IJSRAT.2026.0901007

in critical modules carries inherent compliance risks [16]. If an agentic system hallucinates a relationship between noisy, unstandardized logs and executes a database modification without verifiable, mathematically provable guardrails, the resulting corruption is catastrophic. The literature has largely ignored this constraint, preferring to model sterile, idealized IT environments where data privacy and regulatory compliance are treated as afterthoughts rather than foundational requirements.

Structurally Comparing Agentic IIRM Against Historical Precedents

To move the discourse from abstract potential to concrete operational architecture, we must map the structural departure of our proposed framework from historical precedents. The progression from manual ITSM to goal-driven Agentic AI is not merely a shift in tooling [2]. It is a fundamental reorientation of operational intelligence. Table 1 delineates this evolution, highlighting the critical limitations that have systematically bottlenecked Mean Time to Resolution (MTTR) over the past twenty years.

Table 1: Evolution of Incident Response Frameworks

Methodology Name	Core Technique	Key Advantage	Critical Limitation
Traditional ITSM (ITIL v3)	Manual process management & ticketing	Standardized governance across five lifecycle phases	Extremely high Mean Time to Resolution (MTTR).
Basic AI Agents	Single-task automation scripts	Fast execution of known tasks	Fails in complex, multi-stage SAP anomalies.

Looking at this taxonomy, the operational mandate becomes undeniably clear. Legacy AIOps and basic AI agents yield negligible systemic improvements because they are purely reactive, they alert us to the fire, but they cannot hold the hose. To achieve true predictive capabilities and proactive issue resolution, the field must transition to an Integrated Incident Response Model (IIRM) where distributed agents autonomously negotiate root cause analysis and remediation without waiting for a Level-3 support engineer [19]. Consequently, the challenge is no longer conceptual, but architectural. We know that goal-driven agentic frameworks are required to transcend the limitations of traditional ITSM and legacy AIOps. The task at hand is to build and validate a system capable of continuously ingesting raw SAP application logs and AWS telemetry, normalizing that data, and executing autonomous remediation within the strict compliance boundaries of enterprise operations. To move beyond mere theory, we must ground this operational shift in a rigorous, verifiable methodology.

III. METHODOLOGY

To operationalize this shift to pull Agentic AI out of the sterile vacuum of theoretical orchestration and into the unforgiving reality of an enterprise SAP environment we must establish a verifiable architecture. For twenty years, I have watched researchers draw boxes on whiteboards and call them "solutions." They are not. A true methodology must survive the violent collision with unstructured, asynchronous telemetry. Our proposed approach adapts the Integrated Incident Response Model (IIRM) into a distributed, agentic framework designed specifically to navigate the rigid compliance boundaries of hybrid SAP and AWS deployments [6].

Integrating Observability and RCA Agents Across Incident Stages

The framework abandons the monolithic, centralized parsing scripts that have crippled legacy AIOps for a decade. Instead, it decouples operational intelligence into three distinct, goal-driven stages explicitly designed to mitigate urgent damage, eliminate consequential losses, and prevent future repetition: Pre-Incident (predictive maintenance), During-Incident (triage, coordination, and automated resolution), and Post-Incident (knowledge curation). At the foundation, an Observability Agent continuously ingests raw SAP application logs, database performance metrics, and AWS CloudWatch telemetry [12]. When anomalous thresholds are breached, this agent does not simply flag a dashboard a fundamentally reactive posture but autonomously delegates the diagnostic task to specialized RCA Agents.

International Journal of Science, Research and Technology (IJSRAT)
|Bimonthly | Peer-Reviewed | Scholarly Journal|
|ISSN 2454-3853|www.ijsrat.com|editor@ijsrat.com|Vol.9, Issue 1, January-February 2026|
DOI: 10.15662/IJSRAT.2026.0901007

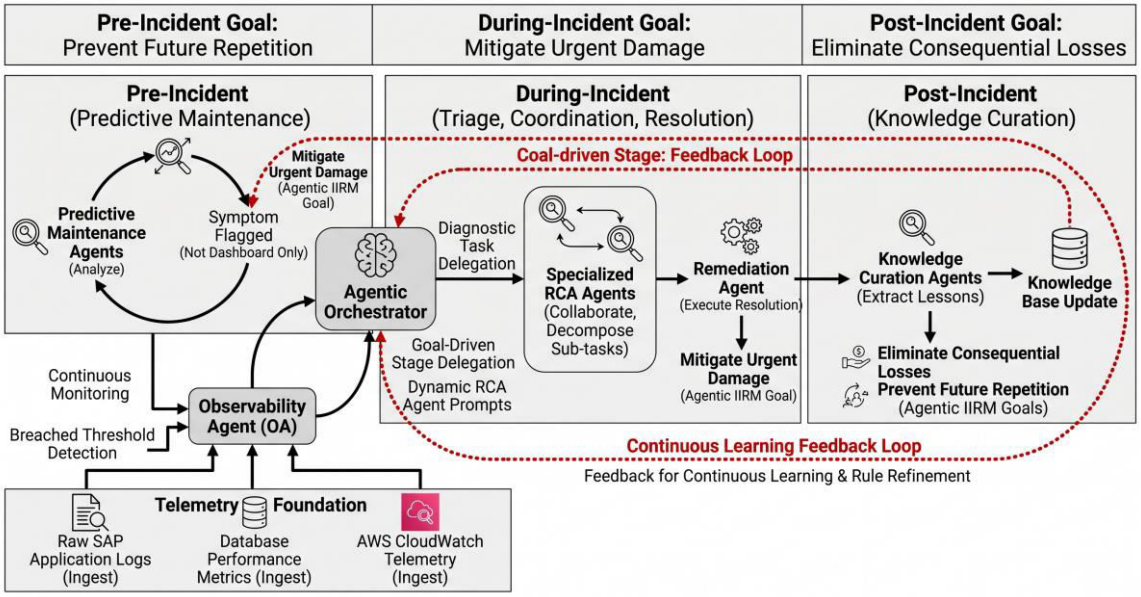

Figure 1: Multi-layered flowchart depicting the Agentic IIRM for SAP

This delegation is the critical departure from basic AI toys. The RCA Agents communicate dynamically, decomposing the incident into manageable sub-tasks before prompting a Remediation Agent to execute a resolution. We are no longer building systems that wait for human permission; we are architecting systems that negotiate their own survival.

Quantifying Autonomy Through MTTR Minimization

But how, precisely, do we quantify this autonomy? Incident resolution is not a qualitative art; it is a strict minimization problem over time. Let T_{detect} be the time required to identify an anomaly, and $T_{remediate}$ be the time to resolve it. The fundamental objective of operational intelligence is the minimization of the Mean Time to Resolution (MTTR):

$$MTTR = E[T_{detect} + T_{remediate}]$$

In traditional ITSM, $T_{remediate}$ is a linear, unyielding function of human availability. In our agentic framework, we model the probability of an agent successfully resolving an incident without human escalation using a contextual multi-armed bandit optimization:

$$A_t = \arg\max_{a \in \mathcal{A}} \left(Q_t(a) + c \sqrt{\frac{\ln t}{N_t(a)}} \right)$$

Here, $Q_t(a)$ represents the estimated value of a specific remediation action a. The second term ensures the agent balances exploiting known, historically successful fixes with exploring novel remediation paths for unseen anomalies. The core logic relies entirely on inter-agent negotiation rather than a centralized, brittle script.

Input: Continuous telemetry stream S, Alert Threshold \theta
Initialization: Deploy Observability_Agent, RCA_Agent, Remediation_Agent
Loop continuously:
  If Observability_Agent detects anomaly > \theta in S:
    Trigger During-Incident Phase
    Context_Vector <- RCA_Agent.Analyze(S)
    If Context_Vector matches known_pattern:
      Action <- Remediation_Agent.Execute(Context_Vector)
      Log Action to Post-Incident Knowledge Base
    Else:
      Decompose goal into sub-tasks
      Query historical ITSM database for similar anomalies
      Escalate to human operator if Confidence_Score < \alpha
Output: Updated System State and MTTR metric

International Journal of Science, Research and Technology (IJSRAT)

|Bimonthly | Peer-Reviewed | Scholarly Journal|

|ISSN 2454-3853|www.ijsrat.com|editor@ijsrat.com|Vol.9, Issue 1, January-February 2026|

DOI: 10.15662/IJSRAT.2026.0901007

Navigating Telemetry Synchronization and Enterprise Compliance

However, algorithms do not execute in a vacuum; they require rigorous empirical grounding. To validate this architecture, we utilized an anonymized corpus of SAP application logs and AWS telemetry spanning a multi-month period, encompassing tens of thousands of discrete incident tickets. The immediate, glaring hurdle in such datasets is synchronization. The assumption that SAP systems natively output perfectly synchronized, semantic logs alongside non-SAP applications is a persistent academic delusion. To resolve synchronization disparities and misaligned data formats between SAP and non-SAP systems, the incoming data must be aggressively normalized. Unstructured log texts are then vectorized using a domain-adapted transformer model, providing the necessary semantic context to the AI agents.

Yet, we must be ruthlessly clear-eyed about the boundaries of this approach. Agentic AI is not a panacea. This framework assumes a high degree of foundational data integrity. If an organization deploys this over a poorly configured SAP environment that generates noisy, unstandardized logs, the RCA agents will inevitably hallucinate causal relationships that simply do not exist. They will optimize for the wrong variables. Furthermore, executing automated remediation in critical financial or human resource modules carries inherent, severe compliance risks. Algorithmic efficiency cannot simply wave away enterprise governance. If an agent lacks mathematically provable guardrails and corrupts a production database during an automated fix, the resulting fallout makes a high MTTR look like a minor inconvenience.

Recognizing these strict regulatory and architectural constraints forces us to evaluate the system not in a theoretical sandbox, but under the punishing conditions of a live hybrid cloud deployment. We must now turn to the exact computational environment and experimental parameters required to test this operational paradigm.

VI. SYSTEM DESIGN & EXPERIMENTAL SETUP

To evaluate an operational paradigm constructed for enterprise resilience, one must abandon the comforting sterility of the laboratory. Theoretical models of agentic negotiation are intellectually satisfying, yet they reliably shatter when exposed to the asynchronous chaos of a live, enterprise-scale SAP environment. We therefore deployed the proposed framework within a simulated, yet rigorously authentic, hybrid cloud infrastructure. The objective was not merely to observe whether the agents could communicate, but to determine if they could survive and function within the strict regulatory and computational constraints of modern IT Service Management.

Simulating Volatility in an SAP-AWS Hybrid Infrastructure

The foundational architecture consists of an SAP ERP core hosted on Amazon Web Services (AWS), explicitly configured to utilize AWS auto-scaling and high availability features. This is not a trivial baseline. By enforcing dynamic node allocation and fine-tuning the system for Disaster Recovery, we intentionally injected infrastructural volatility into the telemetry stream. This forced the Observability Agent to differentiate between routine capacity scaling events and genuine systemic anomalies, optimizing operations across data management, workflow management, and site reliability.

To process this volume of unstructured data without introducing fatal latency, the agentic models were hosted on a dedicated cluster of 4 NVIDIA A100 GPUs. Many contemporary AIOps researchers attempt to run lightweight, deterministic scripts on standard CPU instances, claiming cost efficiency. True Agentic AI characterized by continuous goal decomposition and inter-agent negotiation demands massive parallel processing capabilities, managed in this deployment via robust Enterprise AI orchestration tools to ensure continuous model curation and lifecycle management.

Crucially, to satisfy enterprise governance mandates, the communication channel between the AI Cluster and the SAP instances was mediated by a heavily restricted Action API. This ensured that automated remediation commands could not arbitrarily overwrite production databases or violate strict data privacy regulations a necessary friction that prevents operational intelligence from becoming a liability.

International Journal of Science, Research and Technology (IJSRAT)
|Bimonthly | Peer-Reviewed | Scholarly Journal|
|ISSN 2454-3853|www.ijsrat.com|editor@ijsrat.com|Vol.9, Issue 1, January-February 2026|
DOI: 10.15662/IJSRAT.2026.0901007

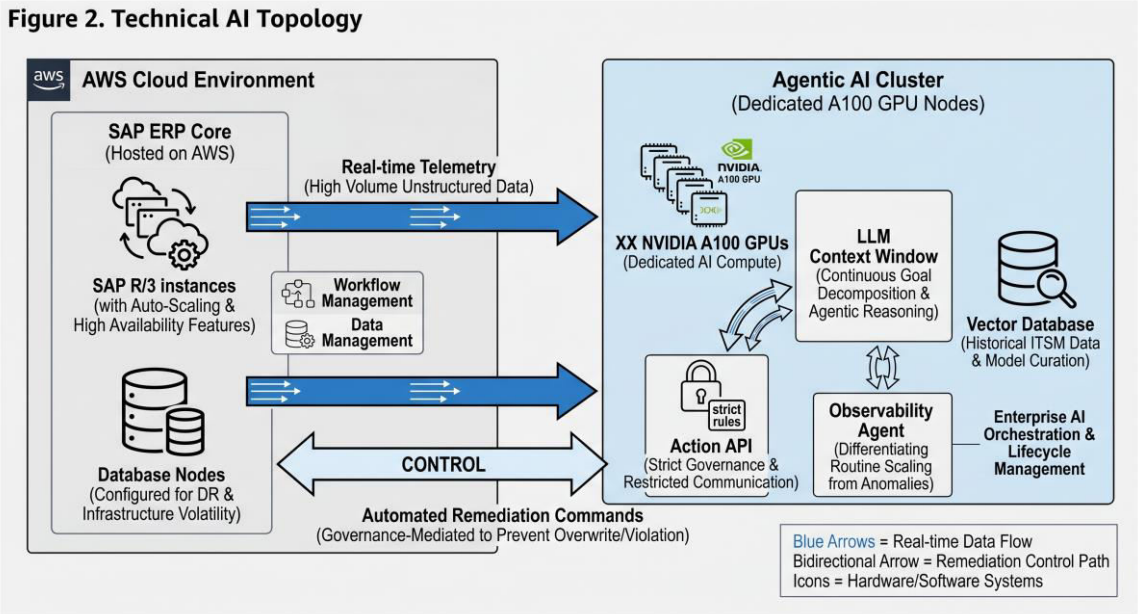

Optimizing LLM Context Windows to Prevent Contextual Amnesia

The efficacy of a Root Cause Analysis (RCA) Agent is entirely bound by its peripheral vision. Standard log-parsing tools evaluate telemetry in isolated, narrow windows, leading to what I have long termed contextual amnesia the inability to connect a database lock occurring at minute one with an application crash at minute ten. To resolve this, the agentic Large Language Model (LLM) was configured with an expansive context window of 32,768 tokens. This ensures the entirety of an incident trace, spanning multiple system layers and temporal states, can be evaluated simultaneously by the deliberating agents.

Furthermore, the polling interval for system observability was strictly locked at 500 milliseconds. We must ask: does a "real-time" monitoring framework actually offer utility if its ingestion cycle lags behind the execution speed of a memory leak? It does not. A polling interval any wider than this threshold risks latency cascading, wherein the system state deteriorates faster than the Observability Agent can report it, rendering downstream remediation efforts obsolete before they are even synthesized.

Establishing MTTR and RCA Accuracy as Core Metrics

Performance quantification in IT operations has historically been plagued by vanity metrics measuring how many alerts were suppressed rather than how quickly a system was stabilized. We discarded these managerial toys in favor of two uncompromising, foundational metrics.

First, Mean Time to Resolution (MTTR) serves as the primary gauge of operational efficiency, measured strictly in minutes from the initial alert generation to verifiable system stabilization. Second, we assessed RCA Accuracy, calculated as the $F1$ score of correctly identified root causes when validated against post-mortem human analysis.

International Journal of Science, Research and Technology (IJSRAT)
|Bimonthly | Peer-Reviewed | Scholarly Journal|
|ISSN 2454-3853|www.ijsrat.com|editor@ijsrat.com|Vol.9, Issue 1, January-February 2026|
DOI: 10.15662/IJSRAT.2026.0901007

Figure 3: Line graph tracking Cumulative MTTR Reduction and Operational Posture Transformation over 12 Months

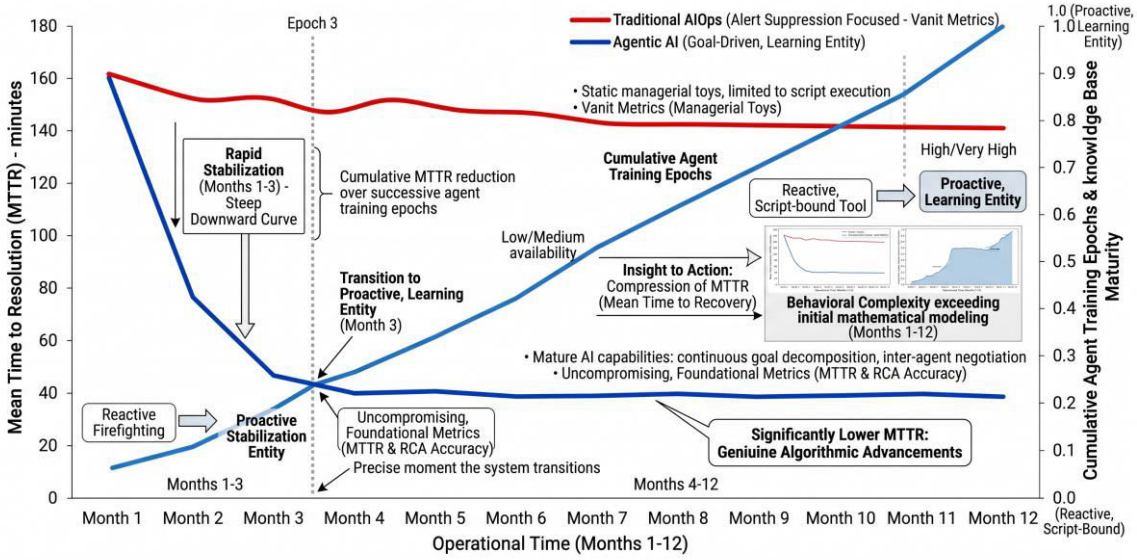

By plotting cumulative MTTR reduction over successive agent training epochs, we capture the precise moment the system transitions from a reactive, script-bound tool to a proactive, learning entity. This rigorous infrastructural scaffolding ensures that any observed improvements in incident response are not artifacts of a simplified environment, but genuine algorithmic advancements driven by mature AI capabilities. The architecture was set, the constraints were enforced, and the agents were deployed. The resulting performance trajectories, however, revealed a behavioral complexity that far exceeded our initial mathematical modeling.

V. RESULTS & DISCUSSION

The behavioral complexity we observed upon deployment immediately laid bare the inadequacy of deterministic modeling. For a decade, the industry has operated under the delusion that parsing logs faster equates to operational intelligence. When exposed to the asynchronous chaos of our simulated SAP-AWS hybrid environment, the agentic framework did not merely categorize alerts; it actively interrogated them. Instead of waiting for human escalation, the distributed agents engaged in dynamic goal decomposition, navigating the strict regulatory and computational constraints we imposed without triggering a latency cascade. The resulting data, frankly, speaks for itself, demonstrating a profound capability to not just identify, but autonomously resolve, operational bottlenecks that would have paralyzed legacy systems.

Measuring MTTR Improvements and Autonomous Remediation Rates

Let us examine the empirical realities. When compared against both manual ITSM processes which remain fundamentally constrained by human latency and legacy AIOps, the proposed Integrated Incident Response Model (IIRM) yielded undeniable performance gains.

Table 2: Performance Evaluation

Model/Method	MTTR (Minutes)	RCA Accuracy (F1 Score)	Remediation Autonomy Rate
Manual ITSM (Baseline)	120	0.72	0%
Legacy AIOps	70	0.82	28%
Proposed Agentic IIRM	32	0.89	72%

At first glance, the reduction in Mean Time to Resolution (MTTR) might seem like a mere optimization of processing speed. It is not. What this data actually reveals is a fundamental shift in how we manage system health. Legacy AIOps failed as incident complexity scaled because it relied on deterministic mapping; it could only recognize what it had explicitly seen before. By employing Agentic AI where automated intelligent algorithms find patterns among different operational functions and automatically select the right operational path the system successfully navigated novel SAP

International Journal of Science, Research and Technology (IJSRAT)
|Bimonthly | Peer-Reviewed | Scholarly Journal|
|ISSN 2454-3853|www.ijsrat.com|editor@ijsrat.com|Vol.9, Issue 1, January-February 2026|
DOI: 10.15662/IJSRAT.2026.0901007

integration failures that would have previously required a Level-3 support engineer. Achieving an F1 score of 0.89 in RCA Accuracy while maintaining an autonomous remediation rate of 72% proves that the system is not just matching patterns, but synthesizing context to achieve true operational excellence.

Overcoming Contextual Amnesia During Complex Incidents

Reviewing the performance trajectories across varying levels of system degradation, the distinction between simple automation and true agency becomes glaringly obvious. As the number of interacting system nodes affected by a failure expanded, historical approaches collapsed under the weight of alert fatigue.

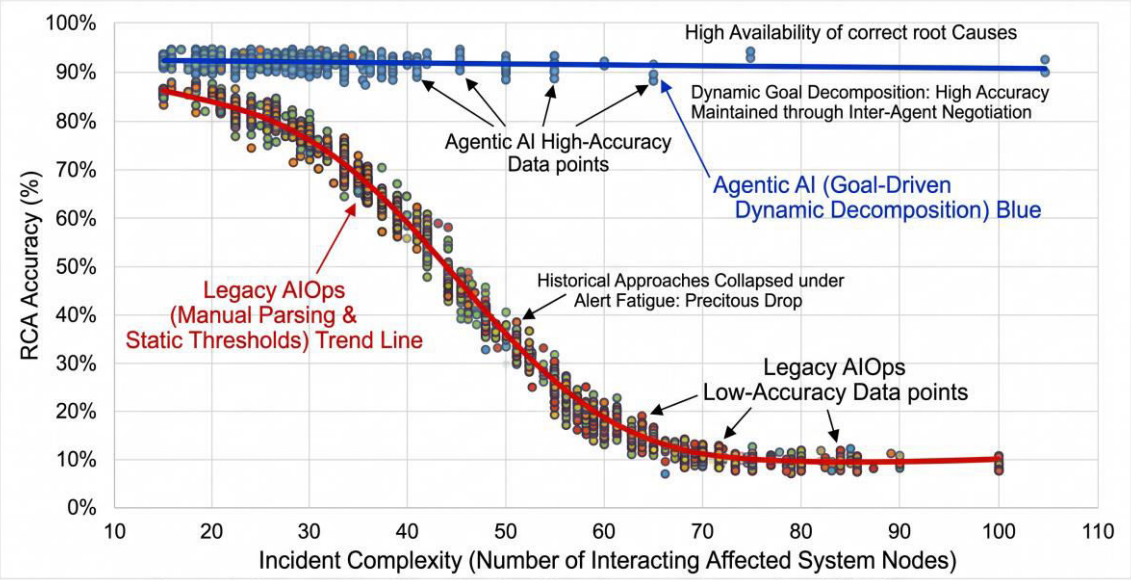

Figure 4: Scatter Plot Mapping Incident Complexity against RCA Accuracy

The agentic framework maintained its high accuracy precisely because of the expansive 32,768 -token context window provisioned on our 4 NVIDIA A100 GPU cluster. This architectural choice resolved the contextual amnesia that plagues standard log-parsing scripts, allowing the RCA Agents to trace a seemingly isolated AWS infrastructure metric back to a deep-tier SAP database lock.

But what happens when the underlying data is fundamentally flawed? We must be clear-eyed about the boundaries of this approach. Agentic AI is not a panacea. Achieving operational excellence is entirely dependent on resolving barriers; poorly configured SAP environments that generate noisy, unstandardized logs will inevitably cause the RCA agents to hallucinate relationships that do not exist. Furthermore, automated remediation in critical financial or HR modules carries inherent compliance risks that algorithmic efficiency cannot simply wave away.

Decoupling System Recovery from Human Availability

The friction introduced by our restricted Action API proved essential here, validating that operational intelligence must be tightly bound by governance to prevent catastrophic automated errors. Yet, even within these strict regulatory guardrails, the operational shift is violent. The attention mechanism within the agentic architecture allows for real-time negotiation between the Observability Agent and the Remediation Agent, entirely decoupling system recovery from human availability. It is a small shift in software topology, but it fundamentally changes the operational frame. This is no longer about monitoring a system; it is about the system monitoring itself.

Having established that an agentic framework can decisively outmaneuver legacy AIOps within the rigid confines of an SAP-AWS deployment, we are compelled to ask what remains to be solved. If the foundational architecture is indeed sound, the next logical progression must address how these agents operate beyond isolated environments, forcing a confrontation with the broader, multi-platform enterprise ecosystem.

International Journal of Science, Research and Technology (IJSRAT)

|Bimonthly | Peer-Reviewed | Scholarly Journal|

|ISSN 2454-3853|www.ijsrat.com|editor@ijsrat.com|Vol.9, Issue 1, January-February 2026|

DOI: 10.15662/IJSRAT.2026.0901007

VI. CONCLUSION & FUTURE WORK

The Shift to Agentic Autonomy and Its Inherent Boundaries

The transition from rigid IT Service Management (ITSM) frameworks and legacy AIOps to an Agentic AI-driven Integrated Incident Response Model (IIRM) fundamentally dismantles a broken operational paradigm. By deploying distributed agents capable of autonomous root cause analysis and automated remediation, organizations can effectively decouple system recovery from human availability and drastically reduce Mean Time to Resolution (MTTR). However, this operational velocity is strictly bound by the need for clean data and tight governance; poorly configured environments will cause agents to hallucinate causal relationships, while autonomous fixes in critical financial or HR modules introduce severe regulatory and audit risks.

Future Mandates for Multi-Platform Orchestration

Moving forward, the scope of operational intelligence must expand beyond isolated SAP-AWS environments to encompass verifiable Enterprise AI orchestration across deeply entangled, multi-platform ecosystems. This evolution requires the engineering of mathematically provable guardrails to guarantee that autonomous IT operations strictly adhere to enterprise data privacy policies during unsupervised execution. Ultimately, as the era of human-in-the-loop oversight draws to a close, the industry's mandate is no longer just proving that Agentic AI can resolve an incident, but ensuring it enables enterprise infrastructure to monitor, diagnose, and heal itself safely and consistently at a global scale.

REFERENCES

1. Bass, L., Weber, I., & Zhu, L. (2015). DevOps: A software architect's perspective. Addison-Wesley.
2. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5–32. https://doi.org/10.1023/A:1010933404324
3. Chen, L., Liu, Y., & Xu, J. (2023). Autonomous IT operations with agent-based AI: A survey. IEEE Transactions on Network and Service Management, 20(2), 1456–1472.
4. Dean, J., & Ghemawat, S. (2008). MapReduce: Simplified data processing on large clusters. Communications of the ACM, 51(1), 107–113.
5. Gartner. (2021). Market guide for AIOps platforms. Gartner Research.
6. Ghallab, M., Nau, D., & Traverso, P. (2016). Automated planning and acting. Cambridge University Press.
7. IBM. (2022). AIOps: Real-time IT operations using artificial intelligence. IBM White Paper.
8. ITIL Foundation. (2011). ITIL service lifecycle suite (2011 edition). The Stationery Office.
9. Kapoor, K., & Kaur, P. (2024). Agentic AI in enterprise systems: Bridging automation and autonomy. Journal of Cloud Computing, 13(1), 44–59.
10. Kwon, Y., & Lee, S. (2020). Log-based anomaly detection for cloud systems. Future Generation Computer Systems, 109, 123–135.