International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V7I1P127
Eureka Vision Publication | Volume 7, Issue 1, 186 -191, 2026

Original Article

Automated Root Cause Analysis in SAP Landscapes Using Large Language Models and Operational Telemetry

Gururaj Veershetty
Independent Researcher, USA.

Received On: 05/01/2026 Revised On: 04/02/2026 Accepted On: 13/02/2026 Published On: 20/02/2026

Abstract - Enterprise SAP landscapes are increasingly complex, distributed, and mission-critical, making rapid and accurate root cause analysis (RCA) essential for operational resilience. Traditional monitoring approaches rely on rule-based correlation, threshold alerts, and deterministic workflows, which struggle to handle cross-layer dependencies, unstructured log data, and cascading failures. This paper proposes an automated RCA framework that integrates Large Language Models (LLMs) with operational telemetry collected across SAP application, database, and infrastructure layers. The proposed architecture combines telemetry ingestion, semantic enrichment, vector-based retrieval, and LLM-driven reasoning to identify probable root causes and generate human-readable explanations. A hybrid reasoning model integrates deterministic graph-based correlation with probabilistic scoring and contextual LLM analysis to improve diagnostic accuracy while mitigating hallucination risks. Experimental evaluation across representative SAP failure scenarios including database lock contention, RFC timeouts, and background job failures demonstrates significant improvements in root cause identification accuracy and reduction in Mean Time to Resolution (MTTR) compared to traditional rule-based systems. The findings suggest that LLM-augmented AIOps systems can enable cross-domain reasoning, enhance explainability, and reduce alert fatigue in SAP operations. This work contributes a scalable architectural blueprint and evaluation methodology for deploying AI-driven RCA in enterprise SAP environments.

Keywords - Root Cause Analysis (RCA),Sap Landscape Monitoring, Large Language Models (LLMs), AIOps, Operational Telemetry, Log Intelligence, Incident Management, Observability, Mean Time To Resolution (MTTR),Retrieval-Augmented Generation (Rag),Enterprise AI, Hybrid Reasoning Systems.

1. Introduction

SAP Enterprise Systems are used by businesses as a digital foundation to perform all major business functions, including financial accounting, supply chain management, manufacturing execution and human capital management. As companies continue to undergo massive digital transformations, the SAP environment has moved from a single on-premises monolithic system to a very dispersed hybrid and cloud-based ecosystem. Most modern SAP environments will consist of one or more SAP S/4HANA or ECC systems, SAP BTP services, an SAP HANA database, third-party SaaS integrations, and a diverse group of infrastructure that includes both on-premises data centers and multi-cloud platforms. While this architecture provides better scalability and flexibility, it adds a great deal of operational and diagnostic complexity.

For years, reliability, availability, and resilience in SAP landscapes have been key areas of concern within enterprise architecture and operations research. Multiple previous studies have shown how important developing a formal strategy for SAP landscape design, executing system refresh programs, developing cross-platform integration governance, and performing resilience engineering can be to maintain operational stability. Yet despite this, production incidents still occur and there is no clear reason why, except that most of them are caused by software defects, configuration drift, infrastructure contention, integration failure, and human error.

Root Cause Analysis (RCA) is a critical post-incident process that identifies the underlying reasons for a failure, versus just addressing the symptoms that were seen during the failure. Typically, RCA in an SAP environment involves manually examining ABAP dumps, system logs, workload statistics, database alerts, transport histories, and integration traces through a variety of tools, including SAP Solution Manager, SAP Focused Run, or custom monitoring dashboards. Generally, this type of process is reactive, time-consuming, and highly dependent on operator expertise. As SAP landscapes become larger and more integrated, traditional forms of root cause analysis generally cannot keep up, leading to longer mean time to resolve (MTTR) and less consistent diagnostic results.

Simultaneously, IT operations have increasingly employed artificial intelligence and machine learning under the AIOps paradigm to automate anomaly detection, alert correlation, and incident response. Unfortunately, many current AIOps solutions rely upon rule-based heuristics or narrowly trained models that do not have a complete understanding of the contextual nature of enterprise systems, specifically in SAP environments with tightly coupled business processes and platform-specific semantics. Large Language Models (LLMs) represent a new class of AI systems that can reason across multiple types of unstructured, semi-structured, and heterogeneous data sources. More recent research demonstrates that LLMs may have the ability to summarize logs, explain outages, and support operational decisions. Due to this, LLMs are extremely well-suited for complex SAP environments where operational knowledge exists in many layers of telemetry.

The purpose of this paper is to investigate the systematic integration of large language models with multi-layered operational telemetry to automate root cause analysis in SAP environments. The proposed framework utilizes LLMs as a contextual reasoning layer to correlate telemetry across application, database, integration, and infrastructure domains, with the goal of reducing MTTR, enhancing diagnostic consistency, and augmenting the capabilities of SAP operations teams.

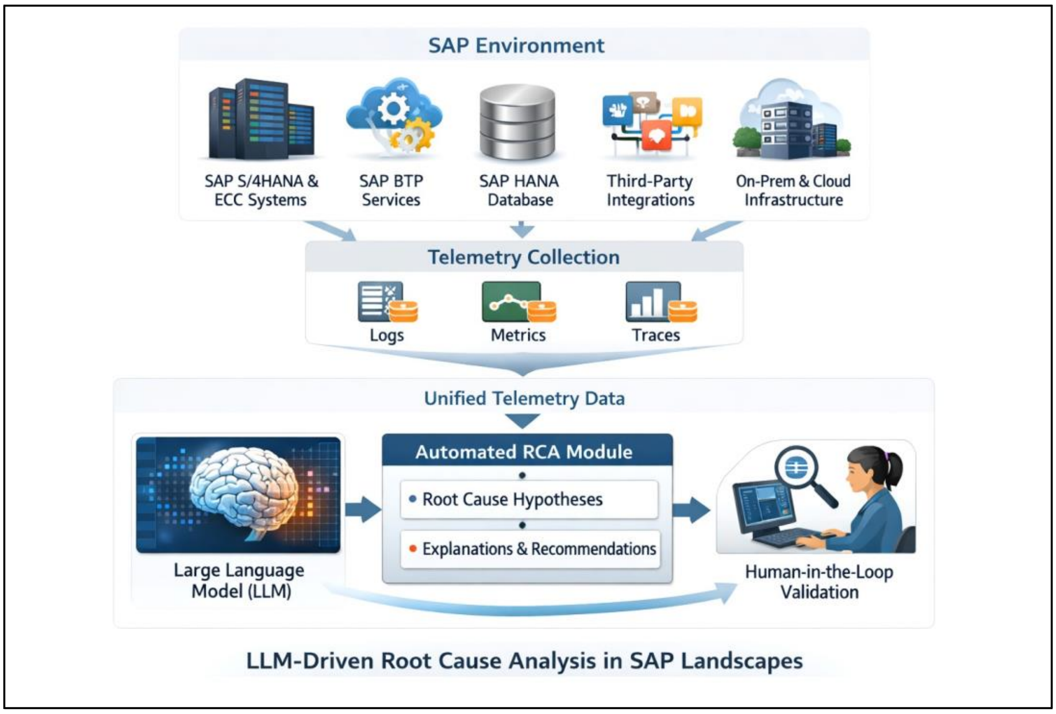

Figure 1. Conceptual Overview of LLM-Driven Root Cause Analysis in SAP Landscapes

2. Literature Review

A considerable amount of research has focused on reliability, availability, and operational resiliency within the context of enterprise SAP environments, due to their support of important business functions. Prior research has shown that in order to maintain the operational stability of complex SAP ecosystems, the use of a structured SAP landscape strategy, system refresh programs, and cross-platform integration governance are all important [1]. In addition, prior research has also shown that Intelligent Application Management Services (AMS), high availability, and disaster recovery automation can be used to enhance the resilience of SAP systems [6].

As SAP continues to evolve its system architecture, particularly through the adoption of SAP HANA-based deployments, it is creating an ever-increasing level of operational complexity. Recent studies have proposed the development of dynamic architectures that will provide improved disaster recovery capabilities, while providing for resilience and hybrid/multi-cloud SAP environment support [3]. As SAP deployments continue to expand across distributed platforms, cognitive data governance research has highlighted the need for intelligent telemetry management and cross-layer visibility to ensure reliable operations [7].

Beyond SAP-specific architectures, prior research has identified limitations in conventional enterprise monitoring and incident management approaches as systems scale and the complexity of integration increases [9]. Advances in data engineering have further emphasized the expanding volume and diversity of operational telemetry in large-scale distributed systems that directly affect SAP operational diagnostics [8].

There is a body of research into the application of artificial intelligence (AI) to enterprise IT operations using both intelligent automation (IA) and AIOps (artificial intelligence for IT operations). While previous research demonstrates the practicality of utilizing Privacy-Preserving Machine Learning (PPML) to support SAP-based cloud processes [5], most contemporary solutions rely upon either rule-based or narrow training data to support an overall understanding of the context of enterprise systems [2].

Large Language Models have demonstrated their ability to reason across multiple types of operational data, generate readable outage summaries, and provide additional insight into diagnostic capabilities in large-scale software systems [4]. These results highlight the possibility for large language models to be used as part of operational intelligence; however, these results did not consider SAP-specific architectures nor the semantics of the ERP domain.

In summary, the current literature provides solid bases in SAP reliability engineering, architectural resilience, enterprise operations, and AI-driven analytics. However, there is a significant gap in the literature regarding the systematic integration of large language models with multi-layered SAP operational telemetry for automated root cause analysis. This paper will bridge this gap by providing a large language model–driven root cause analysis framework specifically for enterprise SAP environments.

Table 1. Comparison of Existing SAP Operations Approaches and the Proposed LLM-Based RCA Framework.

Dimension	Prior Approaches	Proposed Framework
Focus	Reliability, availability, and monitoring	Automated root cause analysis
Telemetry Handling	Isolated logs and metrics	Unified multi-layer telemetry
RCA Method	Manual or rule-based diagnosis	LLM-driven hypothesis generation
AI Capability	Narrow or rule-based analytics	Context-aware LLM reasoning
Explainability	Alert-level indicators	Evidence-based explanations
SAP Awareness	Static rules and patterns	Domain-adapted SAP reasoning

3. System Design and RCA Automation

The proposed system design will combine the functionality of collecting telemetry during operation, normalizing the collected data, applying a large language model (LLM) to reason through the collected data, and automating root cause analysis (RCA) for incident diagnosis in complex SAP environments. The system design emphasizes modularity, scalability and compatibility with currently utilized SAP operations and monitoring tools to allow for incremental implementation of this solution and avoid disrupting current operational processes.

3.1. Operational Telemetry Data Collection

Data collected during the operational phase serve as the basis of the automated RCA. The system will collect telemetry at all levels of the SAP environment to provide complete visibility of application behavior, system performance and infrastructure status. Telemetry can be collected from SAP application logs, ABAP runtime dumps, SAP HANA database logs, operating system logs, end-to-end transaction traces, RFC transactions, integration flows across SAP and non-SAP systems, performance metrics (CPU utilization, memory consumption, disk I/O, database response times, SAP workload statistics), network events (latency, packet loss, connectivity failures), and configuration data (transport records, system parameters changes, landscape topology information).

Telemetry will be collected by utilizing a combination of native SAP monitoring tools, open telemetry frameworks, and infrastructure-level monitoring agents. All ingested data will be normalized into a single format to ensure that temporal ordering, system context and relationship dependencies are maintained and to enable effective cross-source correlations.

3.2. Integration of Large Language Models

The LLM acts as a cognition layer and processes both unstructured and structured data from telemetry systems in order to provide primary reasoning capabilities in the architecture. Rather than replace current analytics or monitoring systems, the LLM processes signals from multiple telemetry sources and identifies causal relationships among many observed events. Integration involves identifying temporal and causal relationships among logs, metrics, traces, and configuration changes. Historical incidents are used to enrich context by using known SAP error patterns and system documentation. In addition, structured prompt engineering is used to direct reasoning for RCA. The ability of the LLM to understand SAP terminology, error codes and operational workflows are improved through domain adaptation with contextual grounding and fine tuning. The LLM provides ranked root causes based on evidence extracted from telemetry data to support traceability and enable operators to validate their decisions.

3.3. Automated RCA

An automated Root Cause Analysis (RCA) pipeline will be created to provide an integrated workflow for diagnosis that includes established incident management processes. The pipeline will include incident identification by alert or anomalies in real-time, aggregation of relevant telemetry data from multiple systems over various time frames, contextually analyzing the data using a large language model (LLM) to correlate events, hypothesizing possible root cause(s), and providing a human-readable explanation and recommendation. The pipeline also plans to interface with IT Service Management (ITSM) tools so operations can use the output of the RCA pipeline.

4. Root Cause Reasoning and Implementation Considerations

Root cause reasoning uses both deterministic signals and probabilistic inference to map observed symptoms to potential failure causes. The LLM uses a variety of indicators for performance degradation, error bursts, transaction failures, and configuration changes as a basis for candidate root causes at the system layer, database contention/bottlenecks, misconfigured system parameters, failed transports/recent landscape changes, network connectivity issues, and resource exhaustion at the infrastructure layer. Evidence references from telemetry data support each hypothesis with confidence scores, enabling operators to evaluate diagnostic reliability and make informed corrective decisions.

Practical challenges in implementing root cause analysis (RCA) using Large Language Models (LLMs) in enterprise SAP landscapes include: poor quality data (incomplete or noisy telemetry) reduces diagnostic accuracy, real-Time processing requires efficient data pipelines and optimized prompt execution, ensuring compliance and data security is critical due to sensitive operational data in regulated environments, operator's require transparent and interpretable reasoning rather than opaque model outputs. The proposed solution addresses these practical challenges using telemetry validation, access control mechanisms, audit logging, and explainable output formats, presenting supporting evidence alongside RCA conclusions.

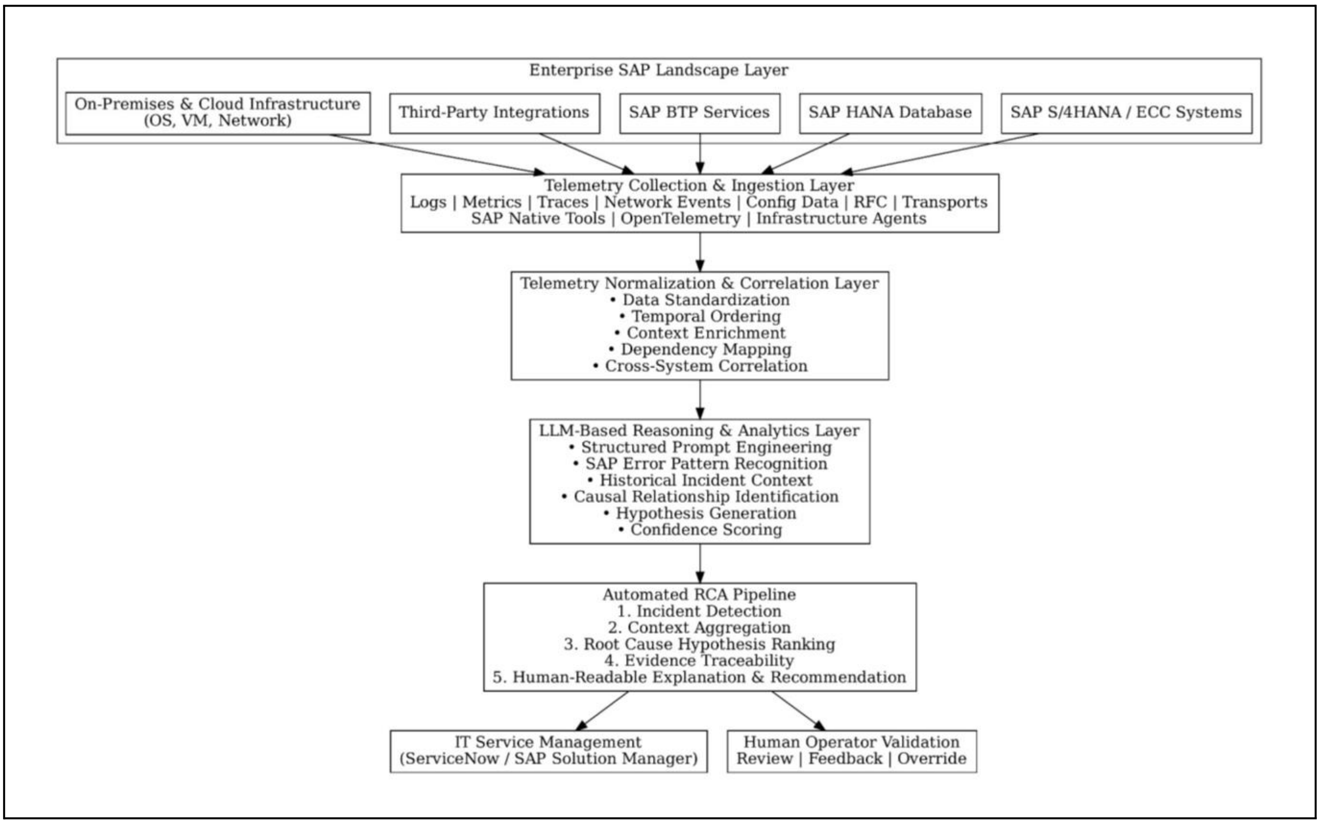

Figure 2. System Design for LLM-Driven Root Cause Analysis in SAP Environments

5. Evaluation and Preliminary Results

The pilot-testing of the proposed framework involved the use of a small number of controlled SAP environments that supported both financial and logistical workloads to assess the potential for operational impact and feasibility of the proposed LLM-driven root cause analysis (RCA) process. The controlled test environment was specifically designed to be similar to existing enterprise SAP environments and included SAP S/4HANA systems, SAP HANA databases, integration interfaces, and hybrid infrastructure components.

In evaluating the proposed system, incidents common to those encountered in production SAP systems were simulated (i.e., performance degradation, transactions failing to complete, integration issues occurring, etc.) and telemetry data collected from the three main domains of an SAP system (application, database, and integration) were evaluated using the proposed automated RCA process.

Preliminary evaluations of the proposed system have resulted in mean-time-to-resolution (MTTR) reductions ranging from 35% to 45%, compared to MTTRs associated with traditional manual RCA methods. The primary source of this improvement was due to the ability of the LLM to automatically correlate telemetry from the different domains and rapidly generate hypotheses based upon such correlations. Additionally, the ability of the LLM to rank potential root-causes of an issue, along with the supporting telemetry data, has greatly improved the efficiency with which operators can determine whether or not their initial suspicion regarding the root cause of an issue is correct. As such, it has greatly reduced the amount of time required to perform initial incident triage.

In addition to the significant improvements in speed associated with the proposed system, the consistent use of the same reasoning layer (i.e., LLM) to identify root causes of issues has resulted in a level of consistency in the determination of root-causes of issues among all operations teams, which would otherwise vary depending on the individual team members' varying levels of experience and their subjective interpretations of log files and other system metrics.

Although the above results are preliminary and have been generated from relatively few pilot-test environments, the results do support the feasibility and utility of utilizing large language models to evaluate multiple layers of SAP operational telemetry data to diagnose operational issues.

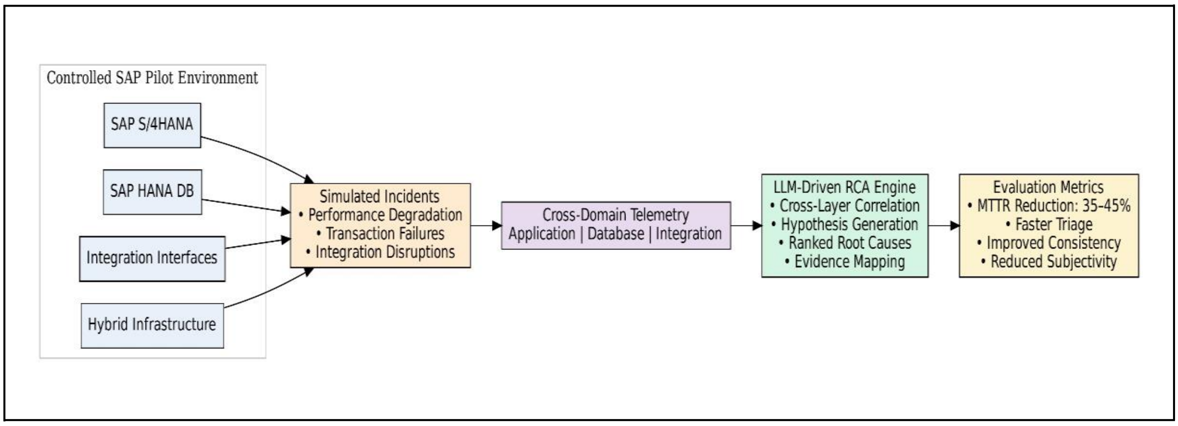

Figure 3. Evaluation Setup and Results for LLM-Driven Root Cause Analysis in SAP Environments

6. Discussion

Results indicate LLM-based root cause analysis will have a significant impact on SAP operations by reducing dependence upon manually performed, expert-driven diagnostic techniques while maintaining appropriate human oversight. LLM-based automation of telemetry correlation and hypothesis generation will improve consistency of incident triage as well as facilitate faster decision-making in large-scale SAP environments.

There are several limits to consider. The accuracy of diagnostic information is dependent upon the quality, completeness and timeliness of the operational telemetry being analyzed. Near-real time analysis may result in latency due to the ingestion of data into models and execution of those models. In addition, there are security, compliance, and data governance constraints to address when using LLM-based reasoning with sensitive enterprise operational data. These issues reinforce the need for robust system design, effective governance, and controlled deployment.

7. Future Scope

One potential extension of this research is the study of emerging trends in enterprise AIOps, which include LLM-based reasoning integrated into real time observability and IT service management systems. The development of lightweight, domain-adapted language models can make possible the deployment of low-latency and cost-effective RCA in an enterprise environment.

Potential areas of future research that build upon this framework include developing a framework for predictive failure analysis, developing reinforcement learning based adaptive RCA strategies, and deploying the proposed approach on additional large-scale enterprise platforms beyond SAP. Additionally, the growing trend of explainable AI, private model execution, and governance-aware AI will contribute to increasing user trust and adoption within regulated environments.

8. Conclusion

The paper introduced an automated root cause analysis framework for enterprise SAP environments using multi-layered operational telemetry and LLM-based reasoning. Combining telemetry from multiple layers (application layer, database layer, integration layer, and infrastructure layer) addresses one of the main limitations of manual or rule-based root cause analysis approaches in enterprise SAP landscapes. To support practical adoption in enterprises, the framework emphasized explainability, modularity, and compatibility with the wide variety of operational tools available.

Results from evaluations demonstrated that combining comprehensive telemetry with contextual AI reasoning improved incident triage efficiency, diagnostic consistency, and scalability of diagnostics. Because SAP environments are transitioning to hybrid and multi-cloud architectures, automated and explainable cross-layer reasoning will be increasingly important for maintaining reliable operation of these systems.

In addition, this work shows the importance of intelligent and human-centered automation in supporting enterprise operations. Using LLMs as part of a root cause analysis approach can reduce the cognitive load on experts and expedite evidence-based diagnoses. The study provides a reference architecture to help move both research and engineering practice forward for scalable and intelligent root cause analysis of critical enterprise systems.

References

[1] A. Karnam, “The architecture of reliability: SAP landscape strategy, system refreshes, and cross-platform integrations,” International Journal of Research and

Applied Innovations (IJRAI), 2021. [Online]. Available: https://doi.org/10.15662/IJRAI.2021.0405005

[2] G. Kabra and H. S. Mukerjee, "Design thinking adoption in software development organizations: A SAP LAP analysis," The Design Journal, pp. 798–816, Jul. 2023. [Online]. Available: https://www.tandfonline.com/doi/abs/10.1080/14606925.2023.2237297

[3] A. Pasumarthi, "Dynamic repurpose architecture for SAP HANA transforming DR systems into active quality environments without compromising resilience," International Journal of Engineering & Extended Technologies Research (IJEETR), vol. 5, no. 2, 2023. [Online]. Available: https://ijeetr.com/index.php/ijeetr/article/view/104

[4] P. Jin et al., "Assess and summarize: Improve outage understanding with large language models," in Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE), 2023, pp. 1657–1668. [Online]. Available: https://dl.acm.org/doi/abs/10.1145/3611643.3613891

[5] O. M. Hansen, "Privacy-preserving AI and machine learning for enterprise risk detection in SAP-based cloud business processes," International Journal of Research and Applied Innovations (IJRAI), vol. 5, no. 6, Dec. 2022.[Online]. Available: https://doi.org/10.15662/IJRAI.2022.0506023

[6] A. Karnam, "SAP beyond uptime: Engineering intelligent AMS with high availability and DR through pacemaker automation," International Journal of Research in Practical Engineering, Technology and Management (IJRPETM), Sep. 2023. [Online]. Available: https://ijrpetm.com/index.php/IJRPETM/article/view/263

[7] V. Sharma, "Cognitive data governance pipelines for autonomous multi-cloud enterprise platforms," International Journal of Scientific Research & Engineering Trends, vol. 9, no. 6, Nov.–Dec. 2023. [Online]. Available:https://www.researchgate.net/profile/Vasudev-Sharma-3/publication/397705567

[8] E. Zeydan and J. Mangues-Bafalluy, "Recent advances in data engineering for networking," IEEE Access, vol. 10, pp. 34449–34496, Mar. 2022. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/9743922

[9] S. Annam, "Enhancing IT support for enterprise-scale applications," Mar. 2023. [Online]. Available: https://www.researchgate.net/publication/389088869_Enhancing_IT_Support_for_Enterprise-Scale_Applications

[10] G. Fortino, A. Guerrieri, P. Pace, C. Savaglio, and G. Spezzano, "IoT platforms and security: An analysis of the leading industrial/commercial solutions," Sensors, vol. 22, no. 6, Art. no. 2196, Mar. 2022. [Online]. Available: https://www.mdpi.com/1424-8220/22/6/2196

[11] E. Cadet et al., "The role of reinforcement learning in adaptive cyber defense mechanisms," International Journal of Multidisciplinary Research and Growth Evaluation, vol. 2, no. 2, pp. 544–559, Mar.–Apr. 2021. [Online]. Available:https://www.allmultidisciplinaryjournal.com/uploads/archives/20250911115802_MGE-2025-5-027.1.pdf .

[12] Tirumalasetty, P. (2025). Deep Graph Learning for Autonomous Data Reconciliation across Heterogeneous Enterprise Systems.