



A Generative AI-Driven Framework for Automated Enterprise Data Documentation and Metadata Management

Mark Henderson¹

Enterprise Data Governance Architect¹

Paul Richardson²

Research Scientist in Natural Language Processing²

Brian Wallace³

Senior Data Integration Architect³

Chaitanya Srinivas⁴

Senior Java Software Developer⁴

Yashwanth kumar⁵

Senior Solutions Architect⁵

Abstract. Enterprise organizations increasingly depend on complex data ecosystems that include databases, data warehouses, data lakes, cloud platforms, applications, and distributed data pipelines, making accurate and up-to-date data documentation and metadata management a significant challenge. Traditional documentation practices often rely on manual processes, pre-defined templates, spreadsheets, and fragmented institutional knowledge, resulting in documentation gaps, inconsistent metadata, limited data traceability, and increased maintenance effort. This research proposes a Generative AI-Driven Framework for Automated Enterprise Data Documentation and Metadata Management that leverages Generative Artificial Intelligence and Large Language Models (LLMs) to automate the extraction, generation, enrichment, classification, validation, and maintenance of enterprise metadata and documentation. The proposed framework integrates automated schema analysis, metadata extraction, data profiling, semantic interpretation, data lineage analysis, natural-language documentation generation, and metadata validation within a unified architecture. Generative AI analyzes structured and semi-structured enterprise data assets to generate contextual descriptions of tables, columns, business entities, relationships, transformations, and lineage information. A governance layer incorporates metadata standards, validation rules, access controls, quality checks, and human-in-the-loop review mechanisms to improve the accuracy, consistency, and reliability of AI-generated documentation. The framework further supports continuous metadata synchronization by integrating with enterprise data pipelines and data cataloging platforms. The proposed approach aims to reduce manual documentation effort, improve metadata completeness and consistency, enhance data discoverability, strengthen data lineage and governance, and support scalable management of enterprise data assets. The effectiveness of the framework can be evaluated using metrics such as metadata completeness, documentation accuracy, semantic consistency, generation quality, processing time, and reduction in manual effort. Overall, the framework demonstrates the potential of Generative AI to transform conventional enterprise data documentation into an intelligent, automated, scalable, and continuously maintained metadata management capability.



Keywords: Generative Artificial Intelligence, Generative AI, Large Language Models (LLMs), Enterprise Data Documentation, Automated Data Documentation, Metadata Management, Metadata Automation, Enterprise Metadata, Metadata Generation, Metadata Enrichment, Metadata Extraction, Metadata Classification, Metadata Validation, Data Cataloging, Enterprise Data Catalog, Data Governance, Data Quality Management, Data Lineage, Data Traceability, Data Discovery, Data Profiling, Schema Analysis, Semantic Data Management, Natural Language Processing (NLP), Natural Language Generation (NLG), AI-Driven Data Management, Intelligent Data Management, Automated Metadata Generation, Context-Aware Documentation, Knowledge Graphs, Enterprise Knowledge Management, Data Asset Management, Data Integration, Cloud Data Platforms, Data Lakes, Data Warehouses, Structured Data, Semi-Structured Data, Data Pipeline Automation, Intelligent Data Cataloging, Metadata Standards, Metadata Quality, Data Governance Automation, Human-in-the-Loop AI, AI-Assisted Documentation, Continuous Metadata Management, Enterprise Data Architecture, Data Observability, Data Transparency, Explainable AI, AI-Based Data Governance, Semantic Metadata Enrichment, Intelligent Enterprise Systems, Digital Transformation.

I. Introduction

The rapid expansion of digital transformation has resulted in organizations generating and managing enormous volumes of data across databases, data warehouses, data lakes, cloud platforms, enterprise applications, analytical systems, and distributed data pipelines. As enterprise data environments become increasingly heterogeneous, organizations face growing challenges in understanding, documenting, governing, and maintaining their data assets. Data documentation provides essential information about the meaning, structure, ownership, relationships, usage, and transformation of enterprise data. Metadata management complements this process by maintaining contextual information that enables organizations to discover, interpret, classify, trace, and govern data assets throughout their lifecycle. However, conventional approaches to documentation and metadata management frequently depend on manual activities, static templates, spreadsheets, technical specifications, and individually maintained repositories. These approaches become difficult to scale when organizations operate thousands of tables, datasets, applications, interfaces, and transformation pipelines.

Traditional documentation practices also suffer from inconsistency and information loss. Data engineers, architects, analysts, and business users may describe the same data asset using different terminology, resulting in semantic inconsistencies across organizational systems. Documentation may become outdated when schemas, transformation logic, business rules, or data pipelines change. In addition, manually documenting newly created data assets requires considerable time and specialized knowledge. The absence of complete documentation can negatively affect data discovery, governance, compliance, analytics, migration, and decision-making. Consequently, enterprises require intelligent mechanisms capable of automatically generating and continuously maintaining high-quality documentation and metadata.

Generative Artificial Intelligence (Generative AI) provides an emerging opportunity to address these limitations by automatically interpreting data structures and generating meaningful natural-language descriptions. Large Language Models (LLMs) can analyze schemas, column names, data types, business terminology, transformation logic, metadata attributes, and contextual information to generate documentation that is understandable to both technical and business users. Unlike conventional rule-based documentation systems, Generative AI can use contextual relationships and semantic information to produce descriptive content dynamically. When integrated with enterprise metadata repositories and data catalog platforms, these capabilities can support automated metadata generation, enrichment, classification, and maintenance.

This research proposes a Generative AI-Driven Framework for Automated Enterprise Data Documentation and Metadata Management. The framework combines automated metadata extraction, schema analysis, data profiling, semantic interpretation, lineage analysis, Generative AI-based documentation generation, metadata validation, and governance mechanisms within an integrated architecture. The framework is designed to automatically transform technical data information into structured and


human-readable documentation while maintaining connections between technical metadata and business context. It also introduces validation and human-in-the-loop mechanisms to reduce inaccurate or unsupported AI-generated information.

The primary objective of the proposed framework is to reduce the manual effort associated with enterprise data documentation while improving metadata completeness, consistency, accessibility, and maintainability. The framework can support various enterprise environments, including relational databases, cloud data warehouses, data lakes, APIs, ETL/ELT pipelines, and analytical platforms. By enabling continuous metadata synchronization and automated documentation updates, the proposed approach can contribute to stronger data governance, improved data discovery, enhanced lineage visibility, and more effective enterprise data management.

II. Background of Enterprise Data Documentation and Metadata Management

Enterprise data documentation involves the systematic recording of information required to understand and use organizational data assets. Such information may include dataset descriptions, table definitions, column meanings, data types, business definitions, relationships, ownership, sensitivity classifications, transformation rules, and usage information. Documentation acts as a communication layer between technical teams and business stakeholders. Developers and data engineers use technical documentation to understand system structures, while analysts and business users rely on business-oriented descriptions to identify appropriate datasets for analytical and operational purposes.

Metadata represents information that describes other data. It provides contextual information about the structure, meaning, origin, ownership, quality, classification, and lifecycle of data assets. Metadata can generally be categorized into technical metadata, business metadata, operational metadata, and governance metadata. Technical metadata describes schemas, tables, columns, data types, keys, indexes, and database structures. Business metadata provides business definitions, domain terminology, ownership, and usage context. Operational metadata includes information about processing schedules, pipeline execution, refresh frequency, and system dependencies. Governance metadata includes classifications, policies, access restrictions, retention requirements, and compliance information.



Effective metadata management enables organizations to establish a common understanding of their data environment. It supports data discovery, impact analysis, lineage tracking, regulatory compliance, data quality management, and enterprise ana-



lytics. However, metadata becomes difficult to manage when organizations operate multiple technologies and platforms. Differences in naming conventions, schema structures, metadata formats, and governance practices can create fragmented metadata environments. Therefore, automated mechanisms are increasingly required to collect and integrate metadata from heterogeneous enterprise systems.

Challenges in Traditional Data Documentation

Manual Documentation Processes

Manual documentation is one of the major limitations of conventional enterprise data management. Data engineers and analysts may need to inspect database schemas, understand business rules, identify relationships, and prepare documentation manually. As data environments grow, this process becomes resource-intensive and difficult to maintain. Manual documentation can also introduce inconsistencies because different individuals may interpret the same data asset differently.

Documentation Obsolescence

Enterprise data environments are continuously changing. New columns may be added, tables may be modified, pipelines may be redesigned, and applications may migrate to cloud platforms. Static documentation may not automatically reflect these changes. Consequently, documentation can become outdated and unreliable, reducing its usefulness for developers, analysts, auditors, and business stakeholders.

Metadata Fragmentation

Enterprise metadata is often distributed across database systems, data catalogs, ETL platforms, cloud services, spreadsheets, application repositories, and governance systems. Fragmented metadata makes it difficult to establish a unified view of organizational data. Users may need to access multiple systems to understand the complete context of a dataset.

Semantic Inconsistency

Different teams may use different terminology for similar business concepts. For example, customer identifiers may be represented as Customer_ID, CustID, Client_Number, or another organization-specific term. Without semantic interpretation, automated systems may treat these fields as unrelated assets even when they represent the same business concept.

Limited Business Context

Traditional technical metadata frequently explains how data is structured but does not adequately explain why the data exists or how it supports business processes. Generating meaningful business context requires understanding relationships between technical structures and organizational terminology.

Role of Generative AI in Data Documentation

Generative AI introduces intelligent language-based capabilities into enterprise data management. Instead of simply extracting metadata fields, Generative AI can interpret available technical and contextual information and convert it into meaningful descriptions. For example, an AI model can analyze a table containing customer transaction information and generate descriptions for tables, columns, relationships, and business purposes based on available metadata and supporting context.

Large Language Models can process natural-language business definitions, technical specifications, schema information, and documentation to generate structured content. This capability allows enterprises to automate repetitive documentation activities while maintaining a consistent documentation style. Generative AI can also enrich existing metadata by identifying semantic relationships, suggesting business terms, generating descriptions, and classifying data assets.

Another important capability is contextual documentation generation. Instead of generating isolated descriptions for individual columns, an AI-driven system can consider relationships between tables, datasets, business concepts, and transformation processes. This creates more meaningful documentation that explains the broader context of an enterprise data asset.

III. Proposed Generative AI-Driven Framework

The proposed framework consists of several integrated layers that collectively support automated enterprise data documentation and metadata management. The architecture begins with enterprise data sources and metadata extraction mechanisms and progresses through data profiling, semantic analysis, Generative AI processing, metadata validation, governance, and publication.



1. Enterprise Data Source Layer

The data source layer contains heterogeneous enterprise data assets, including relational databases, data warehouses, cloud storage platforms, data lakes, APIs, enterprise applications, and analytical systems. The framework is designed to operate across structured and semi-structured environments.

2. Metadata Extraction Layer

The metadata extraction layer automatically collects technical information from connected data sources. Extracted information may include table names, column names, data types, primary keys, foreign keys, constraints, indexes, schemas, relationships, and system ownership information. Automated extraction reduces the dependency on manually maintained metadata.

3. Data Profiling and Analysis Layer

The profiling layer analyzes data characteristics to provide additional context for documentation generation. Profiling may include null-value analysis, uniqueness analysis, value distributions, data patterns, format detection, and statistical characteristics. These insights can help the AI model generate more meaningful and accurate descriptions.

4. Semantic Interpretation Layer

The semantic interpretation layer identifies relationships between technical metadata and business concepts. Natural Language Processing techniques can be applied to column names, descriptions, business glossaries, data dictionaries, and organizational terminology. This layer helps resolve abbreviations and identify related business concepts.

5. Generative AI Layer

The Generative AI layer represents the core intelligence component of the framework. Metadata and contextual information are supplied to an LLM through controlled prompts or structured generation workflows. The model can generate table descriptions, column definitions, business explanations, metadata tags, classifications, and documentation summaries.

6. Metadata Enrichment Layer

The generated documentation is integrated with existing metadata to create enriched metadata records. Enrichment can include business descriptions, synonyms, classifications, domain associations, ownership information, and semantic relationships. The system can also recommend missing metadata attributes for further review.

7. Validation and Quality Assurance Layer

AI-generated documentation should not automatically be treated as authoritative. The validation layer evaluates generated content against available metadata, predefined rules, business glossaries, and quality requirements. Human reviewers can approve, modify, or reject generated documentation. This human-in-the-loop approach improves trust and reduces the risk of unsupported AI-generated information.

8. Governance and Security Layer

The governance layer applies organizational policies to metadata generation and management. It can incorporate access control, sensitive-data classification, retention policies, audit logging, approval workflows, and metadata standards. Security mechanisms are particularly important when enterprise data is processed by AI systems.


9. Data Catalog and Documentation Publication Layer

Validated metadata and documentation can be published to enterprise data catalogs, knowledge repositories, documentation portals, or governance platforms. Users can search and explore generated documentation through centralized interfaces.

Automated Metadata Generation

Automated metadata generation enables organizations to reduce repetitive metadata creation activities. The proposed framework analyzes technical structures and contextual information to automatically generate metadata attributes. For example, a database column named `cust_dob` can be interpreted as a customer date-of-birth attribute based on naming conventions, data patterns, related fields, and business glossary information.

The generated metadata can include technical descriptions, business definitions, data classifications, semantic categories, relationships, and recommended tags. Metadata generation can be performed during initial data onboarding and subsequently repeated when data structures change. This provides a mechanism for maintaining metadata freshness.



AI-Based Enterprise Data Documentation

The framework uses Generative AI to transform metadata into human-readable documentation. Generated documentation can include descriptions of datasets, tables, columns, relationships, data pipelines, transformations, and business usage. Documentation can also be generated at different abstraction levels.

For technical users, the system can provide detailed schema descriptions, dependencies, and transformation information. For business users, it can generate simplified explanations of business entities, metrics, and data usage. This dual perspective can improve communication between technical and non-technical stakeholders.

IV. Metadata Enrichment and Semantic Understanding

Metadata enrichment extends basic technical metadata with contextual and semantic information. Generative AI can identify relationships between metadata elements and organizational business concepts. For example, multiple fields representing customer identifiers across different systems can be associated with a common business concept such as a customer entity.

Semantic enrichment can also support synonym identification, domain classification, business-term association, and relationship discovery. This capability improves data discoverability because users can search for business concepts rather than relying exclusively on technical field names.



Data Lineage and Traceability

Data lineage provides information about the movement and transformation of data across enterprise systems. The proposed framework can integrate lineage information into generated documentation to explain where data originates, how it is transformed, and where it is consumed.

Generative AI can convert complex technical lineage structures into understandable explanations. For example, instead of displaying only a technical pipeline relationship, the system can generate a narrative describing how source customer information is transformed and incorporated into an analytical customer dataset. Such explanations can support impact analysis, debugging, governance, and audit activities.

Human-in-the-Loop Validation

Although Generative AI can automate documentation generation, human validation remains important for enterprise environments. AI models may generate descriptions that are incomplete, ambiguous, or inconsistent with organizational terminology. Therefore, the framework introduces human review mechanisms.

Data stewards, architects, and subject-matter experts can review generated documentation before publication. Approved content can be incorporated into the enterprise metadata repository, while rejected or modified content can be used to improve subsequent generation workflows. This approach combines AI automation with human expertise.

Data Governance and Compliance

Automated documentation and metadata management can strengthen enterprise data governance by improving visibility into organizational data assets. Metadata can capture ownership, classification, lineage, quality information, retention requirements, and access policies.

The framework can support compliance activities by maintaining consistent documentation and audit trails. Automated metadata generation can also help organizations identify sensitive or regulated data assets that require additional governance controls. However, AI-generated classifications should be validated against organizational policies before they are used for compliance decisions.

Implementation Methodology

The proposed framework can be implemented through a phased methodology. The first phase involves identifying enterprise data sources and establishing metadata extraction mechanisms. The second phase focuses on profiling and analyzing data assets. The third phase integrates Generative AI models for documentation and metadata generation. The fourth phase introduces validation and governance mechanisms. The final phase publishes validated metadata to enterprise data catalogs and continuously monitors changes.

A feedback mechanism can be implemented to capture corrections made by data stewards. These corrections can improve prompt templates, retrieval strategies, validation rules, and generation workflows without necessarily requiring model retraining. This approach provides a practical mechanism for continuous improvement.

Evaluation Framework


The effectiveness of the proposed framework can be evaluated using quantitative and qualitative metrics. Metadata completeness can measure the percentage of required metadata attributes successfully generated. Documentation accuracy can evaluate whether generated descriptions correctly represent the corresponding data assets. Semantic consistency can measure alignment between generated terminology and organizational business glossaries.

Other metrics may include documentation generation time, reduction in manual effort, metadata freshness, data discovery efficiency, human approval rate, and documentation coverage. User evaluations can also be conducted with data engineers, analysts, data stewards, and business users to determine whether AI-generated documentation improves usability and understanding.

Benefits of the Proposed Framework

The proposed framework provides several potential benefits for enterprise organizations. First, it can significantly reduce the manual effort required to document large numbers of data assets. Second, automated generation can improve consistency across documentation repositories. Third, semantic enrichment can improve data discovery and understanding. Fourth, continuous metadata synchronization can reduce documentation obsolescence.

The framework can also improve collaboration between technical and business teams by presenting data information in natural-language formats. Furthermore, integration with lineage and governance capabilities can provide organizations with a more comprehensive understanding of their data environment.



Challenges and Limitations

Despite its advantages, Generative AI-based documentation introduces several challenges. AI-generated content may contain inaccurate or unsupported statements if insufficient context is provided. Enterprise data environments may also contain sensitive information that should not be exposed to unauthorized AI systems. Model selection, prompt design, context management, and access control are therefore important considerations.

Another limitation is the dependency on high-quality source metadata. If schemas, business glossaries, and lineage information are incomplete, the quality of generated documentation may also be limited. Consequently, Generative AI should be considered an augmentation mechanism rather than a complete replacement for data stewardship and governance processes.

Future Scope

Future research can investigate domain-specific LLMs designed for enterprise data management and metadata generation. Retrieval-Augmented Generation (RAG) can be incorporated to provide AI models with organization-specific business glossaries, policies, data dictionaries, and historical documentation. Knowledge graphs can also be integrated to represent relationships between data assets, business concepts, systems, and processes.



Further research can explore autonomous metadata agents capable of monitoring data environments, detecting schema changes, generating updated documentation, and initiating approval workflows automatically. Explainable AI techniques can additionally be investigated to provide evidence supporting generated metadata and documentation. These developments could lead toward continuously self-documenting enterprise data ecosystems.

V. Conclusion

The increasing complexity of enterprise data environments has created a strong need for scalable and intelligent approaches to data documentation and metadata management. Traditional manual approaches are difficult to maintain and often result in incomplete, inconsistent, and outdated documentation. The proposed Generative AI-Driven Framework for Automated Enterprise Data Documentation and Metadata Management addresses these challenges by integrating metadata extraction, data profiling, semantic interpretation, Generative AI, metadata enrichment, validation, governance, and data cataloging into a unified architecture. The framework enables organizations to automatically generate contextual documentation while supporting continuous metadata maintenance and human oversight. By combining artificial intelligence with established data governance practices, the proposed approach can improve documentation efficiency, metadata quality, data discoverability, lineage visibility, and enterprise data management. Overall, Generative AI represents a promising foundation for developing intelligent and scalable data documentation capabilities capable of adapting to continuously evolving enterprise data ecosystems.

References

1. Li, G., Zhou, X., & Zhao, X. (2024). LLM for data management. *Proceedings of the VLDB Endowment*, 17(12), 4213–4216. <https://doi.org/10.14778/3685800.3685838>
2. Srikanth CV. A Unified Agentic AI Framework for Self-Adaptive Quality Engineering in Cloud-Native Financial Architectures. *J Artif Intell Mach Learn & Data Sci* 2026 9(1), 3419-3425. DOI: doi.org/10.51219/JAIMLD/Srikanth-chakravarthyvankayala/680
3. BasiReddy, S. R. (2021). Reframing CRM intelligence through knowledge graph-based relationship modeling. *International Journal of Scientific Research & Engineering Trends*, 7(3). Zenodo. <https://doi.org/10.5281/zenodo.18014115>
4. Vollem, S. (2026). Intelligent platform engineering: Leveraging LLM-driven operational automation for autonomous cloud systems. *International Journal of Scientific Research in Science, Engineering and Technology*, 13(1), 529–545. <https://doi.org/10.32628/IJSRSET2613228>
5. Nanchari, N. (2025). The future of IoT in healthcare: Innovations on the horizon. *European Journal of Advances in Engineering and Technology*, 12(6), 54–56. <https://doi.org/10.5281/zenodo.16022695>
6. Yamsani, N. (2026). Multi-agent autonomous governance networks (MAAGN): A scalable framework for self-regulating AI systems in enterprise data ecosystems. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 12(1), 444–460. <https://doi.org/10.32628/CSEIT261213141>
7. Jahnke, N., & Otto, B. (2023). Data catalogs in the enterprise: Applications and integration. *Datenbank-Spektrum*, 23, 89–96. <https://doi.org/10.1007/s13222-023-00445-2>
8. Seetala SR. Real-Time Data Monitoring Using Cloud Observability Tools: Architectures, Techniques and Emerging Practices. *J Artif Intell Mach Learn & Data Sci* 2023 6(4), 3367-3374. DOI: doi.org/10.51219/JAIMLD/srinivasa-rao-seetala/673
9. Srujana, P. (2025). Architecting LLM powered support bots for data engineering operations: Controlled reasoning, evidence grounding, and audit ready incident workflows. *Journal of Scientific and Engineering Research*, 12(11), 211–225. <https://doi.org/10.5281/zenodo.20201082>
10. Natarajan, G. N., Veerapaneni, S. M., Methuku, V., Venkatesan, V., & Kanji, R. K. (2025). Federated AI for surgical robotics: Enhancing precision, privacy, and real-time decision-making in smart healthcare. In *2025 5th International Conference on Intelligent Technologies (CONIT)* (pp. 1–7). IEEE. <https://doi.org/10.1109/CONIT65521.2025.11167543>
11. Ghanta, S. (2025). Engineering productivity at scale: Designing internal developer platforms for cloud-native Java teams. *International Journal of Scientific Research in Science and Technology*, 12(5), 736–748. <https://doi.org/10.32628/IJSRST25126486>



12. Eichler, R., Giebler, C., Gröger, C., Hoos, E., Schwarz, H., & Mitschang, B. (2021). Enterprise-Wide Metadata Management: An Industry Case on the Current State and Challenges. *Business Information Systems*, 1, 269–279. <https://doi.org/10.52825/bis.v1i.47>
13. Menda, J. R. (2025). Hybrid reasoning developer assistants leveraging LLMs for enterprise grade Java and Node financial systems. *AVE Trends in Intelligent Management Letters*, 1(3), 150–164. <https://doi.org/10.64091/ATIML.2025.000164>
14. Teegala, R. (2025). A systems oriented study of LLM aware design patterns in microservices. *Journal of Scientific and Engineering Research*, 12(12), 157–168. <https://doi.org/10.5281/zenodo.19202588>
15. BasiReddy, S. R. (2023). Human-centered automation frameworks for next-generation CRM platforms. *Journal of Scientific and Engineering Research*, 10(1), 120–127. <https://doi.org/10.5281/zenodo.18467397>
16. Abraham, R., Schneider, J., & vom Brocke, J. (2019). Data governance: A conceptual framework, structured review, and research agenda. *International Journal of Information Management*, 49, 424–438. <https://doi.org/10.1016/j.ijinfomgt.2019.07.008>
17. Nagender, Y. (2025). Transforming master data into strategic infrastructure: Governance maturity, intelligent automation, and enterprise integration leadership. *International Journal of Scientific Research in Science, Engineering and Technology*, 12(6), 538–568. <https://doi.org/10.32628/IJSRSET261345>
18. Vollem, S. (2024). Developing autonomous self-healing infrastructure frameworks using predictive monitoring and intelligent automation to strengthen reliability and resilience in distributed computing environments. *International Journal of Scientific Research & Engineering Trends*, 10(5). <https://doi.org/10.5281/zenodo.19208689>
19. Seetala SR. Intelligent Data Validation in Modern Data Platforms: Integrating Statistical Methods and AI for Reliable Machine Learning Pipelines. *J Artif Intell Mach Learn & Data Sci* 2022 5(2), 3359-3366. doi.org/10.51219/JAIMLD/srinivasa-rao-seetala/672
20. Boeckhout, M., Zielhuis, G. A., & Bredenoord, A. L. (2018). The FAIR guiding principles for data stewardship: Fair enough? *European Journal of Human Genetics*, 26, 931–936. <https://doi.org/10.1038/s41431-018-0160-0>
21. Nanchari, N. (2024). Optimizing healthcare costs and ROI through IoT integration: A strategic evaluation. *International Journal of Science, Engineering and Technology*, 12(6). <https://doi.org/10.5281/zenodo.15791028>
22. Ghanta, S. (2023). From observability to understanding: Automated incident triage using large language model reasoning over logs, metrics, and traces. *International Journal of Engineering & Extended Technologies Research*, 5(5), 7242–7249. <https://doi.org/10.15662/IJEETR.2023.0505009>
23. Parepalli, S. (2022). Semantic and reasoning driven approaches to automated error classification in large scale ETL systems. *European Journal of Advances in Engineering and Technology*, 9(11), 151–162. <https://doi.org/10.5281/zenodo.18084352>
24. Menda, J. R. (2019). A comprehensive study on designing regulatory compliant multi-cloud resilience architectures for mission-critical Java-based enterprise systems. *International Journal of Science, Engineering and Technology*, 7(6). <https://doi.org/10.5281/zenodo.18107819>
25. Vankayala, S. C. (2025). Autonomous quality agents: Policy-driven test generation and intelligent orchestration for continuous software assurance. *European Journal of Advances in Engineering and Technology*, 12(1), 35–42. <https://doi.org/10.5281/zenodo.18467296>
26. Jakubik, J., Vössing, M., Kühn, N., Walk, J., & Satzger, G. (2024). Data-centric artificial intelligence. *Business & Information Systems Engineering*, 66, 507–515. <https://doi.org/10.1007/s12599-024-00857-8>
27. BasiReddy, S. R. (2024). Architecting trustworthy and scalable CRM intelligence with LLM-driven integration and zero trust governance. *Journal of Artificial Intelligence, Machine Learning & Data Science*, 3(2), 2988–2993. <https://doi.org/10.51219/JAIMLD/santhosh-reddybasireddy/620>
28. Teegala, R. (2025). A governance oriented study of policy aware generative AI in contact center and customer service workflows. *International Journal of Scientific Research in Science and Technology*, 12(16), 528–543. <https://doi.org/10.32628/IJSRST25121661>
29. Nagender, Y. (2025). AI-guided decision intelligence for autonomous master data management platforms: Enterprise governance practices at Inspire Brands. *International Journal of Scientific Research in Science and Technology (IJSRST)*, 12(9), 264–301. <https://doi.org/10.32628/IJSRST2512940>
30. Malerba, D., & Pasquadibisceglie, V. (2024). Data-centric AI. *Journal of Intelligent Information Systems*, 62, 1493–1502. <https://doi.org/10.1007/s10844-024-00901-9>



31. Devarapalli, S., Vathsavai, V. G., Katkam, V., & Kanji, R. K. (2025). Cloud-native LLMops meets DataOps: A unified framework for high-volume analytical systems. In 2025 International Conference on Computing Technologies & Data Communication (ICCTDC) (pp. 1–8). IEEE. <https://doi.org/10.1109/ICCTDC64446.2025.11158069>
32. Seetala, S. R. (2021). Master data management as a strategic foundation for enterprise consistency: Frameworks, architectures, and governance practices. *International Journal of Computer Technology and Electronics Communication*, 4(1), 3230–3240. <https://doi.org/10.15680/IJCTECE.2021.0401005>
33. Vollem, S. (2023). Artificial intelligence for root cause analysis in cloud-native systems: Techniques, architectures, and research trends. *European Journal of Advances in Engineering and Technology*, 10(9), 120–129. <https://doi.org/10.5281/zenodo.19347481>
34. Ghanta, S. (2022). Engineering resilience in multi-cloud Java microservices: Architectural patterns across AWS and Google Cloud. *International Journal of Scientific Research & Engineering Trends*, 8(3). <https://doi.org/10.5281/zenodo.1808161>
35. Pipino, L. L., Lee, Y. W., & Wang, R. Y. (2002). Data quality assessment. *Communications of the ACM*, 45(4), 211–218. <https://doi.org/10.1145/505248.506010>
36. Menda, J. R. (2019). A distributed identity orchestration framework for secure authentication automation leveraging Keycloak, OAuth 2.0 grant types, and adaptive access policies. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 5(4), 364–381. <https://doi.org/10.32628/CSEIT192144>
37. Nanchari, N. (2024). Challenges in IoT device interoperability in healthcare. *Journal of Artificial Intelligence, Machine Learning and Data Science*, 2(1), 2727–2728. <https://doi.org/10.51219/JAIMLD/nithin-nanchari/575>
38. Parepalli, S. (2021). Hybrid control strategies for efficient scheduling and flow management in ETL pipelines. *International Journal of Scientific Research & Engineering Trends*, 7(3). <https://doi.org/10.5281/zenodo.17896504>
39. BasiReddy, S. R. (2025). A multilayer governance framework for trustworthy AI-augmented CRM ecosystems. *European Journal of Advances in Engineering and Technology*, 12(7), 80–86. <https://doi.org/10.5281/zenodo.17949783>
40. Vankayala, S. C. (2024). Intelligent quality assurance in cloud-native systems: A deep learning and reinforcement learning approach to adaptive test coverage optimization. *International Journal of Scientific Research in Science, Engineering and Technology (IJSRSET)*, 11(6), 546–553. <https://doi.org/10.32628/IJSRSET2512555>
41. Simmhan, Y. L., Plale, B., & Gannon, D. (2005). A survey of data provenance in e-science. *ACM SIGMOD Record*, 34(3), 31–36. <https://doi.org/10.1145/1084805.1084812>
42. Nagender, Y. (2024). LLM-augmented enterprise search and knowledge discovery in master data management systems. *International Journal of Scientific Research & Engineering Trends*, 10(3). <https://doi.org/10.5281/zenodo.19130581>
43. Teegala, R. (2024). Designing auditable architectures for generative AI systems in enterprise environments. *KOS Journal of AIML, Data Science, and Robotics*, 1(1), 1–10. <https://doi.org/10.5281/zenodo.18712484>
44. Surasani, V. R., Kanji, R. K., Kalluri, R. R., Lanka, H. R., & Puram, K. (2025). Improving user experience design through usability evaluation of mobile multimedia apps. In 2025 4th International Conference on Creative Communication and Innovative Technology (ICCIT) (pp. 1–6). IEEE. <https://doi.org/10.1109/ICCIT65724.2025.11167004>
45. Seetala, S. R. (2020). Secure data architecture models for protecting sensitive information in distributed enterprise environments. *International Journal of Science, Engineering and Technology*, 8(3). <https://doi.org/10.5281/zenodo.19219998>
46. Patel, H., Guttula, S., Gupta, N., Hans, S., Sharma Mittal, R., & N., L. (2023). A data-centric AI framework for automating exploratory data analysis and data quality tasks. *ACM Journal of Data and Information Quality*, 15(4), Article 44. <https://doi.org/10.1145/3603709>
47. Ghanta, S. (2021). System-level testing of event-driven microservices using reproducible containerized environments. *International Journal of Science, Engineering and Technology*, 9(6). <https://doi.org/10.5281/zenodo.18084378>
48. Vollem, S. (2024). From deterministic pipelines to intelligent orchestration: A transformer-driven framework for LLM-augmented DevOps automation. *International Journal of Research Publications in Engineering, Technology and Management*, 7(1), 9964–9975. <https://doi.org/10.15662/IJRPETM.2024.0701009>
49. Menda, J. R. (2018). A hybrid log-driven and event-time streaming pipeline: Integrating Kafka Streams with Apache Flink for real-time financial transaction processing. *Journal of Scientific and Engineering Research*, 5(1), 284–292. <https://doi.org/10.5281/zenodo.18084933>
50. Parepalli, S. (2019). Event-driven architectures for real-time analytics feeds in enterprise systems. *Journal of Scientific and Engineering Research*, 6(11), 338–349. <https://doi.org/10.5281/zenodo.20200945>



51. Teegala, R. (2023). Retrieval augmented generation driven operational runbooks for cloud native systems. *European Journal of Advances in Engineering and Technology*, 10(6), 107–119. <https://doi.org/10.5281/zenodo.19565112>
52. Vankayala, S. C. (2023). Reinforcement learning-driven cognitive testing for scalable and resilient financial systems. *ESP Journal of Engineering & Technology Advancements*, 3(4), 209–217. <https://doi.org/10.5281/zenodo.20092735>
53. Nanchari, N. (2023). Real-time health monitoring and alerts via IoT. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology (IJSRCSEIT)*, 9(4), 710–713. <https://doi.org/10.32628/CSEIT23564523>