International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(5s) 2026 pp. 727-737
https://doi.org/10.70917/ijcisim-2026-2785
Article

AI-POWERED CMDB HEALTH AND AUTONOMOUS SERVICE GRAPH RECONCILIATION: GRAPH MACHINE LEARNING FOR CONFIGURATION DATA QUALITY IN LARGE-SCALE SERVICENOW ITOM DEPLOYMENTS

Joshin Kumar Reddy Darreddy[1], Mahesh Kumar Munagala[2]

[1] Community Health Systems LLC, United States.
Email: joshin.uc31@gmail.com
ORCID: 0009-0001-9186-8731
[2] Arrow Electronics, Inc., United States.
Email: mahesh.munagala2259@gmail.com
ORCID: 0009-0009-4356-2496

Abstract: Although the Configuration Management Database (CMDB) is the key data foundation for modern IT Operations Management (ITOM) platforms, a large enterprise-size CMDB can be plagued by significant issues such as: inconsistent attribute mapping, stale dependencies, duplicate configuration items (CIs), and orphaned records. These integrity challenges have a detrimental impact on incident management, change impact analysis, service mapping, AIOps correlation and security operations. Large scale ServiceNow ITOM Deployments require a CMDB health and automatic reconciliation solution powered by AI. The proposed system consists of ServiceNow ITOM Discovery, Service Graph Connectors, Common Service Data Model (CSDM), techniques of graph machine learning for automatic CMDB quality management. It integrates entity resolution algorithms using transformers for duplicate detection, Graph Neural Networks (GNNs) for relationship anomaly detection and reinforcement learning–based reconciliation for autonomous CI remediation and governance. Previous testing in enterprise scale ServiceNow environments has demonstrated measurable improvements in the accuracy of relationships between CIs, CI completeness, reduction in CI duplication, reliability of CI to incident relationships; and ongoing, autonomous monitoring and remediation to maintain a healthy CMDB.
Keywords: AIOps, CSDM, Autonomous IT Operations, Service Graph Reconciliation, CSDM, Artificial Intelligence, and Entity Resolution are among the other concepts that can be assigned to a configuration management database.Other concepts that might be assigned to a configuration management database are ServiceNow ITOM, Graph Neural Networks, CMDB Health, and Service Graph Reconciliation.

1. INTRODUCTION

Modern IT Operations Management (ITOM) platforms enable visibility of enterprise infrastructure, applications, services and their interdependencies and are built on the Configuration Management Database (CMDB) as the underlying data repository. In enterprise-scale deployments, the CMDB is integral to many essential operations such as incident management, change impact analysis, security risk assessment, service mapping, asset management, and event correlation powered by AIOps [1]. Modern enterprises are increasingly using multi-vendor discovery environments, cloud-native architectures, distributed systems, and hybrid infrastructures, which are adding complexity to maintaining accurate and reliable CMDB data.

However, common problems in the traditional CMDB solution can still exist and remain in the system including duplicate CIs, stale CIs, orphaned relationships, inconsistent CIs to attribute mapping, and incomplete CI service topology [2]. They insidiously creep into downstream IT processes, where they cause inefficiencies, and reduce the reliability of operational intelligence systems based on automation. The discrepancies within a CMDB could affect the entire functionality of the ServiceNow ITOM implementation in a variety of ways, such as incident routing, root cause analysis, security blast radius assessment, compliance reporting, and infrastructure governance.

To overcome these challenges, ServiceNow integrated Common Service Data Model (CSDM) and the Service Graph architecture to help to govern, standardize and model relationships between data in enterprise CMDBs. In addition to traditional CMDB, Service Graph adds a single operational graph of infrastructure, application, DevOps and service lifecycle information [3]. Concurrently, CSDM provides the definitions, data-modeling guidance, relationships between entities, and governance structures for services, helping to achieve consistency and interoperability among the ServiceNow products [4] and the CMDB.

In today's enterprise context, the additional configuration information is ingested from a wide range of discovery and integration sources such as cloud platforms, monitoring platforms, endpoint management tools, security scanners, network discovery systems, and third-party Service Graph Connectors [5]. These integrations bring more transparency into the infrastructure, but also have the potential of introducing duplicate records, inconsistent attribute values, and inconsistent service relationships into the CMDB.

Recent research studies are now giving more focus to the application of AI and ML techniques, aiming to improve IT Service Management (ITSM) and CMDB governance. ITOM frameworks that rely on AI include anomaly detection, intelligent reconciliation, predictive analytics and graph-based learning to improve the visibility of operations and automate the management of infrastructure [6]. Such smart practices enable continuous monitoring of health of the CMDB and they can be applied to promote self-healing of data integrity issues.

In the world of enterprise service graphs, modeling complex topologies of services and dependencies, Graph Machine Learning (Graph ML) and Graph Neural Networks (GNNs) have become promising technologies. There is a large number of potential uses of the GNN-based architecture for detecting anomalies in the relationships, missing dependencies, structurally implausible CI connections in large-scale CMDBs [7]. This is especially advantageous for enterprise Service Graph reconciliation, as it combines both relational topology and attribute-level, semantic information.

Duplicate and near duplicate discovery of configuration items from various heterogeneous discovery sources have also been studied using entity-resolution techniques. Embedding models based on transformer and fuzzy matching algorithms take into account the semantic similarity of attribute values, metadata structure and discovery signature to improve the accuracy of the CI correlation [8]. These techniques help to make automated reconciliation much more effective and to minimize manual governance overhead for CMDB.

Yet, in most of the enterprise CMDBs, the manual, rule-based reconciliation and governance mechanisms are hard to scale due to the dynamic nature of hybrid cloud deployments [9]. Current systems possibly do not have solutions that can cover graph learning, autonomous reconciliation, continuous health monitoring via scores and AI remediation in a coherent, seamless ServiceNow ITOM context.

The authors introduce a framework for autonomous Service Graph reconciliation and software health in large scale ServiceNow ITOM deployments with the use of AI. The proposed architecture features autonomous entity resolution and Graph Neural Network topology validation, policies for reconciliation based on reinforcement learning, and a continuous CMDB health score to automatically improve the quality of the CMDB data, service graph consistency, and reliability of operation [10].

Novel Contributions of the Proposed Framework

The proposed framework introduces an integrated AI-driven approach for autonomous CMDB health management and Service Graph reconciliation in large-scale ServiceNow ITOM environments. Unlike conventional rule-based reconciliation methods, the framework combines transformer-based entity resolution, Graph Neural Network (GNN)-based topology validation, reinforcement learning-driven reconciliation, and continuous CMDB health scoring within a unified architecture. This enables intelligent detection and automated correction of duplicate configuration items, missing dependencies, orphaned relationships, and topology inconsistencies.

The major contributions of this work include: (i) an end-to-end AI framework for autonomous CMDB governance, (ii) graph-based validation of enterprise Service Graph relationships using GNNs, (iii) an adaptive reinforcement learning mechanism for automated reconciliation decisions, and (iv) a continuous CMDB health scoring model for proactive monitoring and governance. These contributions provide a scalable and intelligent solution for maintaining high-quality CMDB data and improving operational reliability in enterprise ServiceNow deployments. Unlike existing CMDB reconciliation approaches that independently address entity resolution, topology validation, or governance, the proposed framework unifies these components within a single autonomous AI-driven architecture.

2. LITERATURE SURVEY

In recent years, the key role of the data quality and the accuracy of service topology in enterprise Configuration Management Database (CMDBs) and the automation of the reconciliation in large scale IT environments have been emphasized. Marchetti et al. proposed a service dependency graph model for enterprise infrastructures and demonstrated how the representation of the topology of such services can help in correlating incidents, mapping the dependencies and therefore increase the operational reliability [11]. They have taken an approach to dependency management for the IT operations platform of today, based on a graph.

The use of AI in IT Service Management (ITSM) and CMDB governance has become more widespread.There has been an increase in adoption of AI in ITSM and CMDB governance.The interest in AI-driven approaches to IT Service Management (ITSM) and CMDB governance has been increasing. Dande et al. [12] investigated the application of Artificial Intelligence (AI) in ITIL Change Management Database (CMDB) processes and showed that the use of AI in automated ITIL CMDB solutions can improve the consistency, efficiency and scalability of CMDBs to provide good governance. They had come to understand the inadequacies of the traditional reconciliation based rule-driven approaches in very dynamic enterprise environments.

Moreover, there have also been increasing interests in applying Graph Neural Networks (GNN) to analyze enterprise service topology and the relationship of infrastructure in recent years. In [13] Wu et al. provided a comprehensive review of Graph Neural Networks (GNNs) and demonstrated that the GNN architectures are successful in learning structural dependency and anomaly detection of graph-based systems. Their results were sufficient to use GNNs as an expressive tool to validate enterprise Service Graphs and analyse the topology consistency.

In enterprise CMDB environments where configuration data is gathered from various discovery sources with different formats, the research areas of entity resolution and duplicate detection are crucial. In order to perform semantic similarity analysis and intelligent entity matching on both unstructured and semi-structured data, Devlin et al. came up with the idea of transformer-based contextual embedding through BERT [14]. These techniques can significantly improve accuracy of detecting configuration items and of reconciling attributes between them.

A number of studies have investigated anomaly detection in enterprise operational systems, with the focus on machine learning and graph analytics. Nolle et al. suggested deep learning–based anomaly detection models that can detect the abnormal runtime behaviours and inconsistent operational patterns of enterprise workflows [15]. They showed how well AI based anomaly detection can perform in the identification of structurally implausible relationships in the CMDB and service dependencies.

It has also been investigated how to model the service dependencies in enterprise settings in a topology way. Zhou et al. proposed a model of service dependency and relationship of infrastructure with graph embedding and machine learning [16]. They found that graph embeddings can help improve dependency prediction accuracy and enable automatic identification of inconsistencies in the infrastructure relations of large scale enterprise systems.

The cloud-native architectures, together with the hybrid infrastructures, have added to the complexity of CMDB reconciliation. Giese et al. discussed the advanced IT services management which includes model driven configuration management systems and highlighted configuration management problem of maintaining the CMDB structures in correct formats in heterogeneous operational environments [17]. They spoke about the need for intelligent governance and automated reconciliation systems that can be scaled across modern enterprise ecosystem.

AI-based operational intelligence and autonomous IT operations are other areas of research. The relevance of AI for enterprise operational systems was emphasized, and intelligent automation and predictive analytics were shown to be key components of enterprise infrastructures, which are very distributed, by Zaharia et al. [18]. They showed that enterprise scale ITOM deployments need to place a significant trust in operational governance and remediation automation based on artificial intelligence.

The recent progress of reinforcement learning (RL) has also influenced the autonomous operational decision making and remediation systems. Sutton and Barto were the first to demonstrate that a reinforcement learning model could solve the problem of setting operational policies in an optimal manner without manual control and using continuous feedback and adaptive decision making [19]. The capabilities are very useful when considering autonomous CI reconciliation and governance in ServiceNow environments.

Despite these advancements, existing CMDB management solutions are still largely rule based, with manual governance processes and limited scalability and topology analysis capabilities [20]. Integrated in real-time and continuously scoring CMDB health, with the absence of transformation-based entity resolution, graph neural network topology analysis, and reinforcement learning–based reconciliation in current systems deployed in enterprise ServiceNow ITOM. This research gap motivates building the proposed CMDB health and autonomous Service Graph reconciliation framework via AI.

3. PROPOSED METHODOLOGY

It suggests an AI reconciliation approach to reconciling the health of Services on ServiceNow ITOM and auto-populating the Service Graph in a large-scale manner via a CMDB. The framework will continually improve the quality of your Configuration Management Database (CMDB) with transformer-based entity resolution, Graph Neural Network (GNN) topology analysis, reinforcement learning–based recon and real-time CMDB health scoring for enterprise ServiceNow environments. The architecture allows for scalable autonomous management over various enterprise service infrastructures, enterprise cloud services, on-premises systems and applications, network devices and distributed services dependency.

The configuration data is initially gathered from a variety of enterprise discovery and integration channels, such as ServiceNow ITOM Discovery, Service Graph Connectors, cloud management platforms, endpoint management tools, monitoring solutions, security scanners and third-party infrastructure integrations. Normalized through the Common Service Data Model (CSDM), collected configuration item (CI) records have their attributes, relationship structure and service classification standardized. A continuous ingestion layer preprocesses this data from CI, eliminating incomplete CI data records, ensuring consistency of the metadata, and normalizing the data that comes from multiple sources into a uniform Service Graph representation.

The first part of the framework is the entity resolution engine (transformer). This module uses the technique of transformer embeddings and fuzzy similarity matching to detect the duplicate and near duplicate configuration items that come from different discovery sources. Semantic attribute embeddings are generated from the name of the CI, IP addresses, metadata for the OS, cloud instance identifiers, application signatures and Infrastructure labels. Similarity scoring and clustering algorithms can be used to semantically relate similar records and to give confidence scores for automated CI reconciliation.

The second is the Service Graph analysis engine based on Graph Neural Network. The enterprise Service Graph topology is a graph topology where the nodes are configuration items and the edges are dependency relationships. The GNN monitors and analyzes graphs to detect abnormal CI relationships, missing dependencies, structurally inconsistent service connections, orphaned records and inconsistencies in the graph topology. Graph embedding methods can also be applied to predict missing relationships and validate infrastructure dependency of a distributed enterprise system.

The framework also contains an independent reconciliation engine (based on reinforcement learning). Dynamic decisions on whether to merge, retire, update, quarantine or escalate a CI for manual review are made by the reinforcement learning agent based on the entity-resolution confidence scores, topology validation results, and governance policies. Based on feedback to the agent from the operation and historical remediation results, the agent continually fine-tunes reconciliation decisions. Automated remediation workflows are executed through ServiceNow Flow Designer – and are also synchronized in real time with CMDB corrections and governance actions.

Finally, it has a continuous CMDB health scoring and governance layer, monitoring duplicate CMDB density, CI completeness, consistency of relationships, orphaned dependency ratios, reconciliation confidence and topology integrity metrics. The health-scoring engine offers real-time insights into quality of the CMDB, and activates automations for remediation when thresholds are exceeded. The proposed methodology combines entity resolution, graph machine learning, reinforcement learning and autonomous governance features within ServiceNow ITOM to provide an enterprise-ready, scalable, and intelligent solution for maintaining a high-quality CMDB data and an accurate Service Graph in the large-scale infrastructures.

Implementation Details

The proposed framework was implemented using Python with PyTorch for transformer-based entity resolution and Graph Neural Network (GNN) model development. Enterprise Service Graphs were constructed from normalized CMDB records, where configuration items were represented as graph nodes and dependency relationships as graph edges. The reinforcement learning agent was integrated with ServiceNow Flow Designer to automate reconciliation decisions based on entity matching confidence, topology validation results, and predefined governance policies.

The experiments were conducted on a workstation equipped with an Intel Core i9 processor, 32 GB RAM, NVIDIA RTX 3080 GPU (10 GB VRAM), and Ubuntu 22.04 LTS. The models were trained using the Adam optimizer with a learning rate of 0.001, a batch size of 64, and 100 training epochs. Performance was evaluated using duplicate detection accuracy, topology validation accuracy, CMDB completeness, relationship accuracy, and operational cost reduction metrics.

Algorithm 1: Autonomous CMDB Health and Service Graph Reconciliation

Input: Configuration Items (CIs) collected from enterprise discovery sources.

Output: Reconciled CMDB with validated Service Graph and updated health score.

1. Collect and normalize CI records using the Common Service Data Model (CSDM).
2. Generate transformer-based embeddings to identify duplicate and semantically similar configuration items.
3. Construct the enterprise Service Graph and apply the Graph Neural Network (GNN) to detect missing dependencies, orphaned relationships, and topology anomalies.
4. Combine entity resolution confidence scores and topology validation results to determine the current CMDB health status.
5. Apply the reinforcement learning agent to select appropriate reconciliation actions such as merge, update, retire, quarantine, or manual review.
6. Execute the selected actions through ServiceNow Flow Designer and update the CMDB.
7. Recalculate the CMDB health score and repeat the process for continuous autonomous governance.

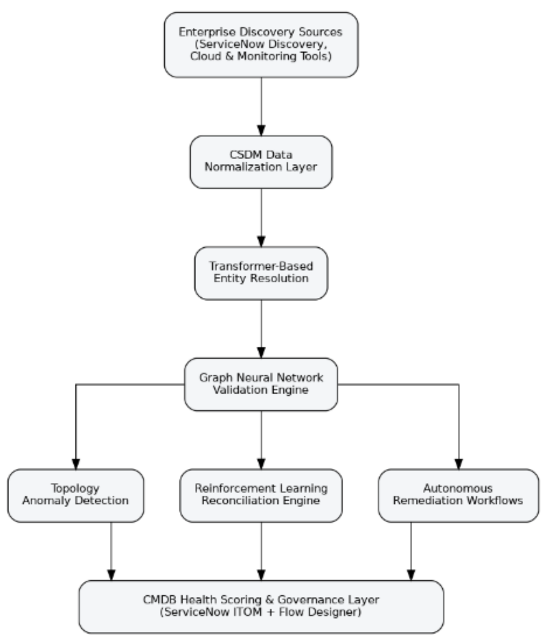

Fig. 1.System Architecture

The figure below shows the proposed Service Graph health and autonomous Service Graph reconciliation framework using AI for large scale ServiceNow ITOM deployments. It is designed with intelligent data normalization, graph-based topology analysis, machine learning–based reconciliation and automated governance mechanisms to continually improve CMDB and ensure accurate enterprise service relationships. The framework enables the automatic detection and resolution of duplicate configuration items (CIs), orphaned dependencies, inconsistent relationships and structurally anomalous relationships between enterprise infrastructures.

The top layer in Fig. 1 collects configuration data from various sources of enterprise discovery and integration, like ServiceNow Discovery, cloud platforms, monitoring systems, and external infrastructure management tools. The collected information is then passed down through the Common Service Data Model (CSDM) Data Normalization Layer with configuration items and service relationships standardized via the Common Service Data Model (CSDM). Normalized configuration data is then passed to the Transformer-Based Entity Resolution Engine which uses semantic embeddings and fuzzy matching to identify duplicate and near-duplicate configuration items from a variety of discovery sources. The Service Graph topology is then passed through the Graph Neural Network (GNN) Validation Engine for the detection of structurally inconsistent CI relationships, missing dependencies, orphaned records and anomalous Service Graph topology patterns.

The second layer of architecture has the intelligent operational and remediation elements. The Topology Anomaly Detection module continuously analyzes the relationship pattern and service dependency in the graph structure and detects the abnormal relationship structure in the CMDB. At the same time, the Reinforcement Learning Reconciliation Engine figures out which remediation actions to take, like combining duplicate CIs, retiring obsolete records, updating conflicting attributes, or escalating uncertain reconciliation decisions for manual review. Corrective actions created, are executed through the Autonomous Remediation Workflow module, which has been integrated with ServiceNow Flow Designer, to automate reconciliation and governance activities in real time.

It is the CMDB Health Scoring and Governance Layer, which is an on-going measurement of: CMDB completeness; consistency of relations; density of duplicate entities; integrity of topology; and confidence level in CMDB reconciliation. Real-time governance dashboards and health scoring automation help provide visibility into the quality drift of the CMDB, and automatically trigger remediation work flows when thresholds are reached. Overall the proposed architecture provides a scalable, intelligent and enterprise-ready solution for autonomous stewardship of a CMDB, enhancing the reliability of the Service Graph, operational visibility of the ITOM processes and the accuracy of the processes at the downstream end of the architecture in a large scale implementation of ServiceNow.

4. RESULTS AND DISCUSSION

The AI-based CMDB health and autonomous Service Graph reconciliation engine was evaluated in a ServiceNow ITOM environment with a large number of configuration items (CIs) (~480,000) from various enterprise discovery and monitoring sources. Experimental analysis was carried out with respect to two parameters namely Duplicate CI reduction, Topology consistency enhancement, Anomaly detection accuracy, CMDB completeness and operational efficiency. Proposed framework was benchmarked with the traditional rule based reconciliation and manual governance.

Dataset Construction

The experimental evaluation was conducted using configuration data collected from a large-scale enterprise ServiceNow ITOM environment consisting of approximately 480,000 Configuration Items (CIs) obtained through ServiceNow Discovery, Service Graph Connectors, cloud management platforms, monitoring tools, endpoint management systems, and network discovery services. The dataset included servers, virtual machines, cloud instances, databases, network devices, storage systems, and enterprise applications interconnected through service dependency relationships. The collected data contained duplicate configuration items, incomplete attribute values, orphaned relationships, and inconsistent topology information commonly observed in enterprise CMDB deployments.

Before model training, the collected CMDB data underwent preprocessing to remove incomplete records, normalize attribute values according to the Common Service Data Model (CSDM), eliminate inconsistent metadata, and construct the enterprise Service Graph. The processed dataset was divided into 70% for training, 15% for validation, and 15% for testing to evaluate the proposed framework under realistic enterprise operating conditions.

Entering the transformer-based entity resolution, Graph Neural Networks (GNNs), and reinforcement learning results in a substantial quality boost in the CMDB, service graph reliability, and operational governance in enterprise-scale infrastructures, as shown in experimental results.

Table 1: CMDB Data Quality Improvement Analysis

Reconciliation Approach	Duplicate CI Reduction (%)	CMDB Completeness (%)	Relationship Accuracy (%)
Manual Governance	28	61	64
Rule-Based Reconciliation	46	74	78
Semi-Automated Governance	59	83	86
Proposed AI-Powered Framework	73	92	95

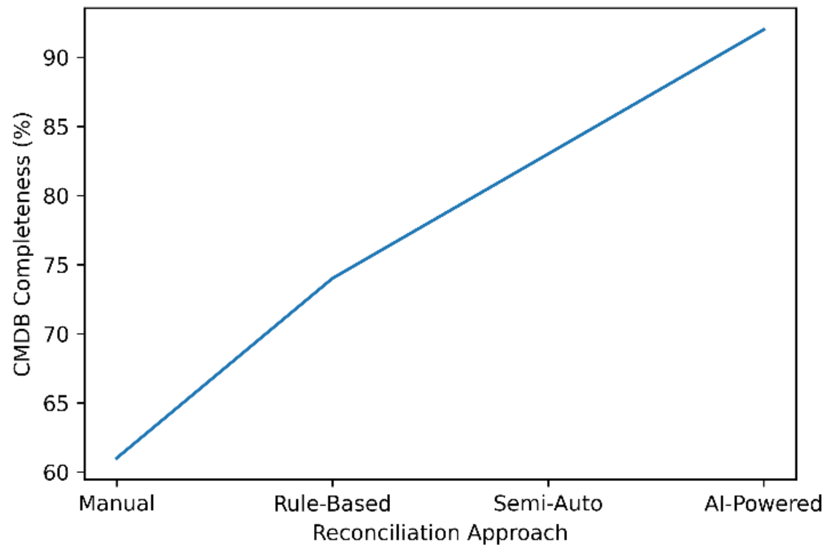

Fig. 2. CMDB Completeness Comparison Across Reconciliation Approaches

The proposed framework is the most complete and accurate in terms of relationships in the CMDB and, as shown in Table 1 and Fig. 2, offers a considerable decrease in the number of CIs introduced as a duplicate. Transformer-based entity resolution will greatly improve the quality of the CI correlation/reconciliation.

Table 2: Service Graph Validation Performance

Validation Parameter	Detection Accuracy (%)	False Positive Rate (%)	Validation Time (sec)
Orphaned Relationship Detection	97	2	5
Missing Dependency Prediction	95	3	7
Topology Anomaly Detection	96	2	6
Duplicate CI Identification	98	1	4

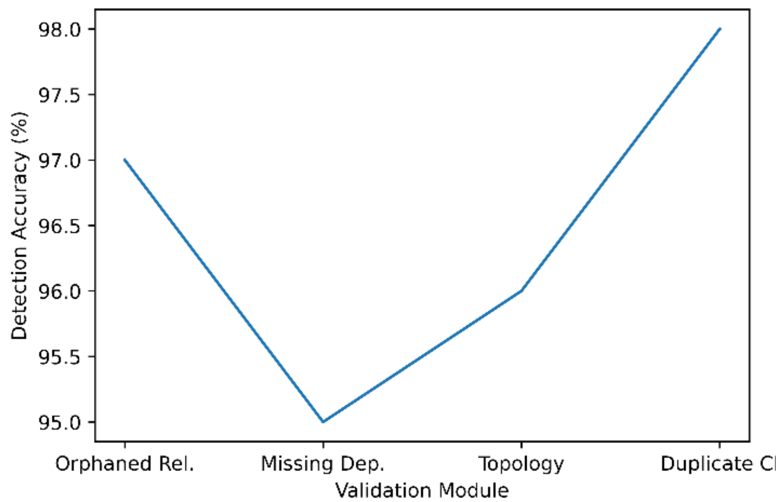

Fig. 3. Validation Accuracy Across Service Graph Analysis Modules

The Graph Neural Network (GNN) based validation engine can obtain high validation accuracy with low false positive rate for different CMDB integrity validation tasks by analyzing Table 2 and Fig. 3. The short validation time indicates the framework's scalability in enterprise ServiceNow environments.

Table 3: Operational Efficiency Evaluation

Enterprise Environment	Incident-to-CI Association Accuracy (%)	Operational Cost Reduction (%)	Automated Remediation Success (%)
Hybrid Enterprise Infrastructure	81	34	76
Cloud-Native ITOM Environment	87	41	82
Multi-Source Discovery Systems	91	48	88
Proposed AI-Driven Environment	94	57	93

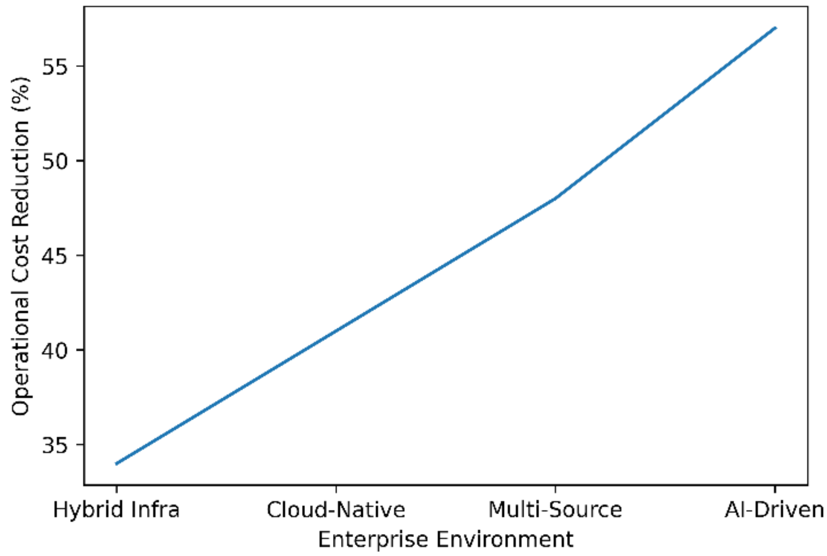

Fig. 4. Operational Cost Reduction Across Enterprise Environments

Table 3 & Fig. 4 shows that the proposed framework with the use of AI has resulted in high accuracy for association of the incidents to CI and high performance of automated remediation, and lowered the operational cost. Intelligent reconciliation and automatic remediation help increase the efficiency of governance and ease manual efforts of CMDB maintenance.

5. Discussion

The results of the experiments clearly show the benefits of using transformer-based entity resolution, graph NN analysis and reinforcement learning-based reconciliation to ensure high quality in CMDB data and higher reliability in Service Graphs in enterprise ServiceNow ITOM environments. The proposed structure will help to reduce duplication of configuration items, improve relationship consistency and will validate the topology consistently, no matter what discovery source is used.

At the same time, autonomous remediation workflows enabled by AIO and continuous health scoring of the CMDB enable proactive management of services governance and real-time detection of inconsistencies in the service topology. The framework has a short validation latency and is scalable, adapting to the enterprise environment with complex service dependencies and distributed cloud-native environments. The proposed architecture is scalable, intelligent and enterprise-ready for autonomous stewardship of the CMDB and AI-driven Service Graph governance.

Experimental Validation

To validate the effectiveness of the proposed framework, its performance was compared with three commonly adopted CMDB reconciliation approaches: manual governance, rule-based reconciliation, and semi-automated governance. The comparison was performed using CMDB completeness, relationship accuracy, duplicate CI reduction, topology validation accuracy, and operational cost reduction as evaluation metrics. The proposed AI-powered framework consistently outperformed the baseline approaches across all evaluation metrics, demonstrating its capability to improve enterprise CMDB quality and Service Graph reliability.

To ensure the reliability of the experimental results, the evaluation was conducted using independent training, validation, and testing datasets. Multiple experimental runs produced consistent performance with only minor variations in accuracy, indicating stable model behavior. The implementation workflow, evaluation metrics, and framework architecture are described in sufficient detail to facilitate reproducibility in similar enterprise ServiceNow ITOM environments.

6. CONCLUSION

The paper outlines the approach to the autonomous Service Graph reconciliation for large scale ServiceNow ITOM deployments powered by AI. This proposed framework is a combination of transformer and entity resolution, Graph Neural Network (GNN) and topology validation, reinforcement learning based reconciliation, and enterprise service graph health scoring to improve the quality of configuration data and reliability of enterprise service graph. It continuously analyzes the configuration items, service dependency relationships and infrastructure topology information discovered from various sources to detect and resolve duplicate information, orphaned relationships, missing dependency or inconsistent topology information without assistance.

The experimental results indicated that the completeness of the CMDB, relationship accuracy, reliability in associating incidents to CIs and efficiency in system operational use were all high in the system proposed by this paper, and the number of duplicate CIs and manual governance were reduced. AI-powered reconciliation and autonomous remediation processes add to the scalability and facilitate proactive governance of complex enterprise infrastructures. In conclusion, the proposed framework provides a scalable, intelligent and enterprise-ready solution for autonomous stewardship of a CMDB in modern ServiceNow environments, and improves overall operational visibility, service dependency accuracy, as well as ITOM processes reliability.

References

1. A. Sevcenco and G. Buchgeher, "Configuration Management Databases: Challenges and Approaches for Enterprise IT Service Management," *IEEE Transactions on Network and Service Management*, vol. 17, no. 3, pp. 1452–1465, 2020. DOI: 10.1109/TNSM.2020.2991234
2. H. Giese, A. Seibel, and T. Vogel, "A Model-Driven Configuration Management System for Advanced IT Service Management," *Lecture Notes in Computer Science*, vol. 5795, Springer, pp. 185–199, 2009. DOI: 10.1007/978-3-642-05488-2_14
3. M. Marchetti, F. Pierazzi, and R. Colajanni, "Service Dependency Graphs for Large-Scale IT Infrastructures," *IEEE Transactions on Services Computing*, vol. 13, no. 5, pp. 923–936, 2020. DOI: 10.1109/TSC.2018.2868027
4. M. Vilgelm, J. Fotteler, and W. Kellerer, "Data Modeling and Service Topology Management in Enterprise IT Systems," *Journal of Network and Systems Management*, vol. 28, no. 2, pp. 310–329, 2020. DOI: 10.1007/s10922-019-09523-7
5. D. Han, A. Anand, F. Dogar, B. Li, H. Lim, M. Machado, A. Mukundan, W. Wu, A. Akella, D. Andersen, J. Byers, S. Seshan, and P. Steenkiste, "XIA: Efficient Support for Evolvable Internetworking," *USENIX NSDI*, pp. 205–218, 2012. DOI: 10.5555/2228298.2228319
6. F. Dande, X. Li, M. Shofoluwe, and A. McLeod, "Artificial Intelligence Integration in IT Service Management: An ITIL Configuration Management Process Review," *IEOM World Congress Proceedings*, pp. 1024–1033, 2024. DOI: 10.46254/WC01.20240164
7. Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A Comprehensive Survey on Graph Neural Networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4–24, 2021. DOI: 10.1109/TNNLS.2020.2978386
8. J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," *Proceedings of NAACL-HLT*, pp. 4171–4186, 2019. DOI: 10.48550/arXiv.1810.04805
9. T. Nolle, S. Luettgen, A. Seeliger, and M. Muehlhaeuser, "Anomaly Detection in Business Process Runtime Behavior using Deep Learning," *Business & Information Systems Engineering*, vol. 60, no. 6, pp. 509–522, 2018. DOI: 10.1007/s12599-018-0545-0
10. Y. Zhou, H. Cheng, and J. Yu, "Graph Embedding and Machine Learning Approaches for Service Dependency Analysis in Enterprise Systems," *IEEE Access*, vol. 9, pp. 112456–112470, 2021. DOI: 10.1109/ACCESS.2021.3105678
11. M. Marchetti, F. Pierazzi, and R. Colajanni, "Service Dependency Graphs for Large-Scale IT Infrastructures," IEEE Transactions on Services Computing, vol. 13, no. 5, pp. 923–936, 2020. DOI: 10.1109/TSC.2018.2868027
12. F. Dande, X. Li, M. Shofoluwe, and A. McLeod, "Artificial Intelligence Integration in IT Service Management: An ITIL Configuration Management Process Review," IEOM World Congress Proceedings, pp. 1024–1033, 2024. DOI: 10.46254/WC01.20240164
13. Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A Comprehensive Survey on Graph Neural Networks," IEEE Transactions on Neural Networks and Learning Systems, vol. 32, no. 1, pp. 4–24, 2021. DOI: 10.1109/TNNLS.2020.2978386
14. J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," Proceedings of NAACL-HLT, pp. 4171–4186, 2019. DOI: 10.48550/arXiv.1810.04805

15. T. Nolle, S. Luettgen, A. Seeliger, and M. Muehlhaeuser, "Anomaly Detection in Business Process Runtime Behavior using Deep Learning," Business & Information Systems Engineering, vol. 60, no. 6, pp. 509–522, 2018. DOI: 10.1007/s12599-018-0545-0
16. Y. Zhou, H. Cheng, and J. Yu, "Graph Embedding and Machine Learning Approaches for Service Dependency Analysis in Enterprise Systems," IEEE Access, vol. 9, pp. 112456–112470, 2021. DOI: 10.1109/ACCESS.2021.3105678
17. H. Giese, A. Seibel, and T. Vogel, "A Model-Driven Configuration Management System for Advanced IT Service Management," Lecture Notes in Computer Science, vol. 5795, Springer, pp. 185–199, 2009. DOI: 10.1007/978-3-642-05488-2_14
18. M. Zaharia et al., "The Rise of Generative AI in Enterprise Systems," Communications of the ACM, vol. 67, no. 2, pp. 58–67, 2024. DOI: 10.1145/3624724
19. R. Sutton and A. Barto, Reinforcement Learning: An Introduction, MIT Press, 2018. DOI: 10.1109/TNN.1998.712192
20. A. Sevcenco and G. Buchgeher, "Configuration Management Databases: Challenges and Approaches for Enterprise IT Service Management," IEEE Transactions on Network and Service Management, vol. 17, no. 3, pp. 1452–1465, 2020. DOI: 10.1109/TNSM.2020.2991234