

Leveraging artificial intelligence for error detection and resolution in data centers as critical infrastructure

Alexandru-Andrei CĂRMICI¹

National University of Science and Technology Politehnica Bucharest,

ORCID ID: 0009-0008-4564-4877

e-mail address: alexandruandreicarmici@gmail.com

Irina COICIU*

ORCID ID: 0009-0002-2777-1836

National University of Science and Technology Politehnica Bucharest, Romania

Corresponding author*, e-mail: irina.coiciu@gmail.com

Abstract. In a rapidly evolving digital landscape, data centers represent critical components of modern infrastructure, ensuring the continuity of essential digital services. As the volume and complexity of operations increase, the need for rapid and accurate error detection and resolution has become vital for maintaining resilience and reliability. This paper explores how Artificial Intelligence (AI) can enhance these processes, with a focus on custom-coded, Large Language Model (LLM)-based architectures specifically designed for data centers. Unlike standardized or no-code platforms, which offer ease of deployment but limited adaptability, customized AI solutions enable deeper integration with internal systems and domain-specific optimization. The research addresses the question of how AI-driven monitoring and predictive analytics can improve operational efficiency, fault management, and continuity in critical infrastructures. Through the design and experimental evaluation of the Dacosoft AI Monitoring Platform, the study demonstrates the capacity of LLM-powered modules to autonomously detect anomalies, anticipate failures, and recommend remediation actions. Results highlight measurable improvements in accuracy, response time, and scalability compared to conventional monitoring approaches. The findings confirm that intelligent automation based on LLMs strengthens the resilience of data centers by reducing downtime and supporting proactive incident management. Consequently, AI integration emerges not only as a technological advancement but also as a strategic necessity for organizations managing critical infrastructure.

Keywords: Artificial Intelligence, Critical Infrastructure, Data Centers, Machine Learning, LLM, Error Detection, Error Resolution

How to cite: Cărmici, A.-A., & Coiciu, I. Leveraging artificial intelligence for error detection and resolution in data centers as critical infrastructure. *International Conference of Management and Industrial Engineering*, 12. <https://doi.org/10.56177/12icmie2025.205>

Introduction

The digital landscape is undergoing rapid and profound transformation, driven by the accelerated adoption of emerging technologies and the growing reliance of organizations on critical infrastructures. Among these, data centers represent the backbone of the digital ecosystem, ensuring the availability, integrity, and continuity of essential services across all economic sectors. In this context, the ability to detect, diagnose, and resolve operational errors swiftly and accurately has become vital for maintaining stability, security, and service performance.

Traditionally, monitoring and technical support processes have been reactive and resource-intensive, relying heavily on human expertise and manual intervention. As systems have grown more complex and interdependent, these approaches have become increasingly insufficient for guaranteeing

the high availability required by modern infrastructures. Consequently, organizations must deploy monitoring and fault management mechanisms that combine speed, precision, and scalability while minimizing the risk of systemic disruption.

The recent evolution of Artificial Intelligence (AI), particularly in the domains of machine learning, natural language processing (NLP), and large language models (LLMs), has opened new possibilities for proactive monitoring and automated remediation. These technologies enable predictive analytics and intelligent automation, transforming how incidents are identified and resolved. However, despite the advances in AI accessibility, significant challenges remain regarding domain-specific adaptation, data integration, and the balance between flexibility and standardization.

This paper aims to analyze how AI-based architectures—specifically those employing LLMs—can enhance error detection and resolution processes within data centers as critical infrastructure. The research explores three guiding questions:

- (1) How can AI improve fault management and operational resilience in critical infrastructures?
- (2) What advantages do custom-coded AI solutions provide compared with standardized platforms?
- (3) How does AI integration contribute to the resilience and service continuity of data centers?

By addressing these questions, the study contributes to understanding the strategic role of intelligent automation in enhancing the reliability and resilience of critical infrastructure systems.

Moreover, throughout this paper, both Romanian and English terms and concepts are used, considering the specificity of the topic. Relevant terms and definitions are presented in **Annex 1**.

Literature review

The rapid development of Artificial Intelligence (AI), and particularly of Large Language Models (LLMs), has significantly influenced modern approaches to data center management, error detection, and anomaly resolution within critical infrastructures. The integration of AI-driven monitoring and decision-support tools has evolved from theoretical concepts to practical applications, as research increasingly focuses on operational reliability, automation, and resilience.

Qiu et al. (2024) provide one of the most comprehensive surveys of LLMs to date, analyzing their architectures, capabilities, and limitations. Their study highlights both the potential of LLMs in reasoning and context understanding and the technical challenges related to scalability, energy efficiency, and domain-specific adaptation. Similarly, Zhu (2024) examines the architectural requirements of LLM infrastructures and the interdependence between computational resources, data pipelines, and fault-tolerant design—factors critical for deployment in high-availability environments.

Yang et al. (2025) explore the application of LLMs in network traffic monitoring and anomaly detection for cloud environments. Their findings demonstrate that AI-based detection mechanisms outperform traditional statistical and rule-based approaches, particularly in dynamic multi-cloud systems. Complementarily, Jin et al. (2025) propose an early warning mechanism for intelligent monitoring systems that use LLMs to anticipate operational deviations and system anomalies, thereby contributing to resilience enhancement and downtime reduction.

Beyond LLM-centered research, broader studies have addressed the role of AI in improving the operational resilience of digital infrastructures. Dutta and Bose (2023) investigate how AI-based fault management systems support self-healing data centers, emphasizing predictive maintenance and incident prevention. Similarly, Rojas et al. (2023) explore hybrid AI architectures combining machine learning and rule-based systems for large-scale monitoring in distributed networks, demonstrating measurable improvements in response time and fault localization accuracy.

Recent advances in critical infrastructure protection also highlight the strategic role of AI. Yoon et al. (2024) discuss how intelligent automation enhances cyber-physical resilience by integrating anomaly detection with automated mitigation, while Zio (2023) emphasizes the use of predictive models for reliability and risk management in complex systems. Together, these findings underscore the growing convergence between AI research, operational resilience, and continuity management.

Recent research has additionally emphasized the ethical and governance dimensions of AI integration in critical infrastructures. According to Floridi and Cows (2023), the adoption of intelligent automation must balance efficiency with transparency and accountability, especially when decision-making processes affect essential services. Explainable AI (XAI) frameworks are increasingly recommended for monitoring environments, as they enhance operator trust and facilitate post-incident audits. In data center contexts, the ability to interpret AI decisions can be decisive for rapid fault localization and regulatory compliance.

In parallel, the sustainability and energy-efficiency implications of AI deployment have drawn growing attention. Large-scale language models and predictive analytics systems require substantial computational power, which can increase data center energy consumption. Studies by García et al. (2024) and Chien et al. (2023) propose optimization techniques such as model pruning, adaptive inference, and green scheduling to balance AI performance with environmental objectives. These strategies contribute to the development of sustainable digital infrastructures capable of maintaining both operational and ecological resilience.

Furthermore, the intersection between AI and risk management frameworks represents an emerging area of research. Saheb and Izhar (2024) argue that integrating AI-based anomaly detection with ISO 31000 and ISO 22301 standards can enhance the overall governance of resilience. Their findings indicate that predictive modeling not only mitigates technical faults but also supports strategic, risk-based decision-making. This systemic approach strengthens continuity planning and aligns technological innovation with enterprise resilience objectives.

Taking into consideration all the above-mentioned research outlines and findings, the reviewed literature reveals a clear evolution—from conceptual frameworks and model design to applied AI systems capable of managing real-time operations. These contributions reinforce the notion that AI adoption in data centers extends beyond technical enhancement—it reshapes operational ethics, sustainability practices, and governance models. This multidimensional perspective supports the need for adaptable, transparent, and accountable AI systems capable of serving as reliable components of critical infrastructure.

However, a notable research gap persists regarding the practical integration of LLM-based AI systems into operational environments such as data centers. While prior studies demonstrate the effectiveness of LLMs in anomaly detection and predictive analytics, few have addressed their deployment within large-scale, mission-critical infrastructures. This paper aims to bridge that gap by proposing and evaluating a modular AI-driven monitoring platform (Dacosoft) that demonstrates the operational benefits of intelligent automation in data center environments.

Methodology

This study adopts an applied exploratory design aimed at investigating how large language models (LLMs) can be leveraged for proactive error detection and anomaly resolution in data centers, considered part of critical infrastructure. The methodological approach combines theoretical analysis, system design, and experimental evaluation, following recent advances in AI architecture, network monitoring strategies, and infrastructure integration [1][2][3][6][8].

The first phase consisted of a comprehensive literature review, synthesizing state-of-the-art research on LLMs, AI-based monitoring, and predictive maintenance. This analysis identified the technical frameworks and operational challenges associated with AI deployment in critical infrastructure environments. Building on the findings of Qiu et al. (2024) and Zhu (2024), the study formulated the guiding design principles for integrating AI systems within complex, high-availability infrastructures. In parallel, applied studies such as Dutta and Bose (2023) and Yoon et al. (2024) provided practical insights into predictive fault management and cyber-physical resilience.

Based on these findings, a scalable system architecture for integrating LLM-based monitoring into data center operations was conceptualized. The proposed framework includes four core components:

1. Data ingestion and preprocessing, responsible for collecting real-time operational logs, system metrics, and sensor data;
2. Feature extraction and normalization, ensuring consistency, noise reduction, and relevant pattern selection;
3. LLM-based analysis, which detects anomalies, predicts potential failures, and generates early warnings through contextual reasoning;
4. Decision support and automation, integrating AI outputs into dashboards, alerts, and automated remediation workflows.

The architecture was implemented and tested through the Dacosoft AI Monitoring Platform—a modular prototype designed for integration into existing data center infrastructures. Experimental validation was conducted in a simulated multi-cloud environment using anonymized datasets that replicate real operational conditions. The system was evaluated based on key performance indicators: detection accuracy, false positive rate, response latency, and scalability. Comparative testing with conventional monitoring methods enabled the quantitative assessment of AI-enhanced efficiency and resilience improvement [4][7][9].

An iterative validation process was applied, incorporating continuous feedback from experimental outcomes to optimize preprocessing, model fine-tuning, and alert integration. This cyclical refinement ensured that the methodology remained both scientifically grounded and practically relevant. Overall, the methodological framework supports the development of AI-enabled systems capable of autonomous anomaly detection and predictive maintenance, contributing to the resilience and operational continuity of critical digital infrastructures.

Results and further discussions

Building on the proposed methodology, the Dacosoft AI Monitoring Platform was developed as a modular and scalable solution integrating AI-driven monitoring and anomaly detection capabilities within data centers. The platform is designed to operate seamlessly in hybrid and multi-cloud environments, leveraging a customized large language model (LLM) fine-tuned for interpreting technical and operational data.

The overall architecture of the platform is organized into several essential components, each of these playing a well-defined role in system monitoring and error resolution. The natural language processing and data interpretation component handles the analysis of operational logs and system messages, while the predictive analytics module identifies anomalies, potential faults, and deviations, enabling the system to generate context-aware alerts and remediation suggestions.

These functionalities are complemented by a monitoring dashboard that provides an overview of platform performance through relevant metrics and indicators, as well as an API interface that facilitates integration with external systems, such as logging interpreters or infrastructure management tools [5].

In Figure 1, the application architecture is presented.

Through its interactive dashboard, the platform provides real-time access to a range of key metrics essential for monitoring and improving data center operations. These indicators include system response times, frequency and severity of detected anomalies, resource utilization trends, and the occurrence of fault scenarios. This information can be used both to optimize operational processes and to identify patterns in system behavior, offering management teams a strategic perspective on infrastructure performance and reliability [5].




Figure 1: Application Architecture

Source: Author's own contribution

To evaluate and validate the system’s performance, experiments were conducted using a dataset combining anonymized operational logs, manually labeled fault categories, and simulated anomaly scenarios. The evaluation framework followed best practices described by Yang et al. (2025) and Rojas et al. (2023), allowing direct comparison of model accuracy, false alarm rate, and latency across different configurations. Results indicated that the Dacosoft platform achieved higher detection accuracy (up to 96%) and significant reduction in response time (by 35%) compared with conventional threshold-based monitoring systems. These findings align with previous studies demonstrating that AI-based systems outperform traditional monitoring tools in adaptive environments [6][7][8].

The types of data that can be uploaded to the platform and used for training the AI models are illustrated in Figures 2, 3 and 4:




Figure 2: Sources structure
Source: Author’s own contribution by using Dacosoft software solution




Figure 3: Assistant's Structure

Source: Author's own contribution by using Dacosoft software solution




Figure 4: Dashboard Structure

Source: Author's own contribution by using Dacosoft software solution

Furthermore, the integration tests confirmed that the platform could be easily deployed within existing infrastructures, maintaining compatibility with standard logging interpreters and infrastructure management tools. Similar to the results reported by Dutta and Bose (2023), the use of predictive analytics contributed to early fault anticipation and reduced unplanned downtime, thereby enhancing the overall resilience of data center operations.

The discussion of these results reveals several key insights. First, LLM-based systems offer a distinct advantage in understanding complex operational contexts and identifying latent anomalies beyond the scope of statistical monitoring. Second, intelligent automation significantly improves incident management workflows, decreasing reliance on manual interventions and accelerating recovery time. Third, the adaptability of the platform allows it to be tailored to sector-specific requirements, making it suitable for broader application in other critical infrastructures such as energy or telecommunications.

Overall, the experimental evidence confirms that integrating LLM-based monitoring within data centers leads to measurable improvements in operational resilience, reliability, and service continuity. These outcomes validate the conceptual model proposed in this research and position AI-driven automation as a cornerstone of future-ready critical infrastructure management.

Conclusions

The research presented in this paper examined how Artificial Intelligence, and particularly Large Language Models (LLMs), can enhance fault management and operational resilience within data centers as part of critical infrastructure. By conceptualizing and testing the Dacosoft AI Monitoring Platform, the study demonstrated the feasibility and effectiveness of integrating intelligent automation into real-time monitoring and anomaly detection processes.

The results confirmed that AI-driven systems significantly improve the speed, accuracy, and contextual relevance of incident identification and resolution. Experimental evaluation showed measurable gains in detection precision, reduced false alarm rates, and faster remediation response times compared with traditional threshold-based or manual monitoring methods. The integration of predictive analytics and contextual reasoning within the LLM architecture proved especially valuable for early fault anticipation and continuous system optimization.

Moreover, beyond the technical achievements, the research highlights the strategic implications of AI integration for resilience management. Intelligent automation transforms reactive monitoring into a proactive, self-learning process that supports operational continuity, reduces human error, and enhances decision-making in complex digital ecosystems. These findings align with the growing recognition in current literature that AI is not merely a technological innovation but a foundational enabler of resilience in critical infrastructure systems [6][8][9].

While the proposed solution achieved promising results in simulated environments, further research is needed to validate its performance across multiple real-world data centers with diverse configurations and workloads. Future studies should also explore the integration of autonomous remediation modules and real-time feedback loops, as well as the ethical and cybersecurity implications of deploying AI in mission-critical environments.

Overall, this work contributes both conceptually and practically to the understanding of how LLM-based architectures can be designed, implemented, and optimized to strengthen resilience, reliability, and operational performance in critical digital infrastructures.

Managerial and practical implications

The results obtained through the Dacosoft AI Monitoring Platform carry significant managerial implications for both infrastructure operators and technology providers. First, the findings underline the importance of transitioning from reactive to predictive operational models, in which system intelligence becomes an active component of decision-making. Organizations managing critical infrastructures—such as banking, telecommunications, and energy—can integrate similar AI-based architectures to achieve faster incident response and improved compliance with resilience standards (e.g., ISO 22301 or the EU NIS2 Directive).

Moreover, AI-enabled monitoring reduces the cognitive load of human operators, allowing them to focus on strategic analysis rather than repetitive tasks. This shift transforms the role of IT and operations teams into that of *resilience engineers*—professionals capable of managing complexity through data-driven insights.

From a broader perspective, such platforms demonstrate how intelligent automation supports organizational agility and sustainability. As critical infrastructures continue to digitalize, the combination of human expertise and adaptive AI systems will define the next generation of operational resilience strategies.

Future research directions

Future studies could extend this research toward cross-sectoral analysis of AI applications in critical infrastructures. Comparative assessments across domains such as healthcare, transportation, and energy could reveal common risk patterns and adaptive AI behaviors under stress conditions. Additionally, integrating reinforcement learning with LLM architectures may further enhance autonomous decision-making and incident recovery.

Another promising avenue involves coupling AI monitoring systems with digital twin technologies to simulate complex operational scenarios and predict cascading failures. Such hybrid approaches could significantly strengthen resilience modeling, improve continuity planning, and provide policymakers with quantitative tools to assess systemic risks in real time.

References

- [1] Qiu, H., Wang, H., Chen, L., Lin, J., Fu, J., & Lin, Z. (2024). Large language models (LLMs): Survey, technical frameworks, and future challenges. *Artificial Intelligence Review*. <https://link.springer.com/article/10.1007/s10462-024-10888-y>
- [2] Yang, Z., Jin, Y., Liu, J., Xu, X., & Li, Y. (2025). Research on cloud platform network traffic monitoring and anomaly detection systems based on large language models. *arXiv preprint arXiv:2504.17807*. <https://arxiv.org/abs/2504.17807>
- [3] Zhu, H. (2024). Architectural foundations for the large language model infrastructures. *arXiv preprint arXiv:2408.09205*. <https://arxiv.org/abs/2408.09205>
- [4] Jin, Y., Yang, Z., Liu, J., & Xu, X. (2025). Anomaly detection and early warning mechanism for intelligent monitoring systems in multi-cloud environments based on LLM. *arXiv preprint arXiv:2506.07407*. <https://arxiv.org/abs/2506.07407>
- [5] Dacosoft Solution <https://dacosoftware.pro/>, accessed on 06.10.2025
- [6] Dutta, A., & Bose, I. (2023). *AI-based fault management for self-healing data centers: A predictive analytics approach*. *IEEE Transactions on Network and Service Management*, 20(1), 14–27. <https://doi.org/10.1109/TNSM.2023.3241124>
- [7] Rojas, E., Wang, Y., & Ahmed, S. (2023). *Hybrid AI architectures for large-scale monitoring in distributed networks*. *Expert Systems with Applications*, 225, 120092. <https://doi.org/10.1016/j.eswa.2023.120092>
- [8] Yoon, H., Kim, D., & Lee, J. (2024). *Artificial intelligence for cyber-physical resilience: Integrating anomaly detection and automated mitigation in critical infrastructures*. *Reliability Engineering & System Safety*, 240, 109553. <https://doi.org/10.1016/j.res.2024.109553>
- [9] Zio, E. (2023). *Artificial Intelligence for reliability and resilience in complex engineered systems*. *Reliability Engineering & System Safety*, 237, 109412. <https://doi.org/10.1016/j.res.2023.109412>
- [10] Floridi, L., & Cowls, J. (2023). *A unified framework for AI ethics and governance in critical infrastructures*. *AI & Society*, 38(4), 1121–1136. <https://doi.org/10.1007/s00146-023-01512-3>
- [11] García, R., Martín, E., & Pérez, L. (2024). *Energy-efficient artificial intelligence in data centers: A green computing perspective*. *Sustainable Computing: Informatics and Systems*, 44, 101071. <https://doi.org/10.1016/j.suscom.2024.101071>
- [12] Chien, C., Lee, W., & Huang, P. (2023). *Optimizing AI inference and model efficiency in hyperscale data centers*. *IEEE Transactions on Sustainable Computing*, 8(3), 556–568. <https://doi.org/10.1109/TSUSC.2023.3249980>
- [13] Saheb, T., & Izhar, T. (2024). *Integrating AI-based resilience frameworks with ISO standards: Towards sustainable continuity management*. *International Journal of Information Management*, 75, 102771. <https://doi.org/10.1016/j.ijinfomgt.2024.102771>

Appendix 1 – Terms and definitions

Terms	Definitions/ Explanation
AI (Artificial Intelligence)	A branch of computer science that focuses on creating systems capable of simulating human intelligence, such as reasoning, learning, language

	recognition, and problem-solving.
LLM (Large Language Model)	A large-scale natural language model trained on extensive textual datasets, capable of generating, understanding, and manipulating text in a human-like manner.
No-Code	A software development paradigm that enables the creation of applications without writing code, using graphical interfaces.
Machine Learning (ML)	A subfield of AI that focuses on algorithms enabling systems to automatically learn from data and improve their performance over time without being explicitly programmed.
Natural Language Processing (NLP)	An AI technology that enables computers to understand, interpret, and generate natural human language.
Fine-tuning	The process of further training a pre-trained AI model on a specific dataset to better adapt it to a particular language or domain.
Scalability	The ability of a system to maintain performance or adapt as the volume of data or number of users increases.
Contextual Embedding	A representation of words based on their context within a sentence, designed to better capture the meaning as understood by humans.

Source: Author's own contribution