The American Journal of Engineering and Technology

ISSN 2689-0984 | Open Access

TYPE Original Research
PAGE NO. 186-193

Check for updates

Resilient Data-Driven Infrastructure: Integrating Predictive Analytics, Secure Pipelines, and High-Performance Fault Diagnosis for Modern Computational Systems

Dr. Elena Morozova
University of Lisbon

OPEN ACCESS

SUBMITED 01 July 2025
ACCEPTED 15 July 2025
PUBLISHED 31 July 2025
VOLUME Vol.07 Issue 07 2025

CITATION

Dr. Elena Morozova. (2025). Resilient Data-Driven Infrastructure: Integrating Predictive Analytics, Secure Pipelines, and High-Performance Fault Diagnosis for Modern Computational Systems. The American Journal of Engineering and Technology, 7(07), 186–193. Retrieved from https://theamericanjournals.com/index.php/tajet/article/view/7066

COPYRIGHT
© 2025 Original content from this work may be used under the terms of the creative commons attributes 4.0 License.

Abstract

Background: Modern computational infrastructures—spanning cloud platforms, GPU-accelerated data centers, and enterprise data lakes—face interlocking challenges: increasing failure rates under intensive workloads, the necessity for secure continuous integration/continuous deployment (CI/CD) practices, and the demand to convert vast heterogeneous data into actionable intelligence (Zhang, 2022; Liu et al., 2023). A coherent, interdisciplinary framework that links predictive analytics, DevSecOps practices, and high-performance fault diagnosis is essential to raise reliability while maintaining scalability and security (Kumar, 2019; Konneru, 2021).

Methods: This article synthesizes theoretical foundations and applied methodologies from the provided literature to produce an integrative conceptual and operational framework. We perform a detailed cross-domain synthesis of techniques from predictive analytics, data engineering and lakehouse architectures, DevSecOps security integrations (SAST/DAST/SCA), high-performance geospatial and GPU computing approaches, and contemporary fault-prediction studies. The methodology includes comparative evaluation of algorithmic families, pipeline architectures, and failure-detection strategies, mapped onto practical system boundaries and operational constraints for cloud and on-premise GPU deployments (Kukreja & Zburivsky, 2021; Li, 2020; Liu et al., 2023).

Results: The synthesis highlights three convergent design principles: (1) unified telemetry and data curation through lakehouse principles to enable low-latency, high-fidelity feature generation (Kukreja & Zburivsky, 2021); (2) embedding predictive analytics into DevOps cycles to create anticipatory operations—thereby improving decision latency and reducing mean time to repair (Kumar, 2019); and (3) a layered fault-diagnosis approach combining supervised predictive models for imminent hardware faults with unsupervised anomaly detection to capture novel failure modes (Peterson et al., 2022; Xie et al., 2021; Liu et al., 2023). The integrated model demonstrates conceptual pathways to reduce unscheduled downtime, decrease non-revenue-impacting water in infrastructure analogues, and enhance security posture inside CI/CD (Kwikima et al., 2024; Konneru, 2021).

Conclusion: A resilient data-driven infrastructure must combine lakehouse data engineering, DevSecOps-integrated CI/CD, and robust predictive diagnostics tailored for high-performance workloads. The proposed framework provides a guide for engineering organizations to structure telemetry, model development, and deployment while accounting for security, scale, and the unique failure characteristics of GPUs and cloud components. Research and industrial practice will benefit from empirical validation campaigns, standardized telemetry schemas, and community-driven benchmarks for fault prediction and remediation orchestration (Liu et al., 2023; Lin & Gupta, 2021).

Keywords: Predictive analytics, DevSecOps, Lakehouse, GPU fault prediction, CI/CD security, High-performance computing, Fault diagnosis

INTRODUCTION

Modern digital infrastructure forms the backbone of a broad spectrum of economic, scientific, and social activity. Cloud platforms and GPU-accelerated computing provide the computational horsepower necessary for machine learning, scientific simulation, and large-scale data services (Zhang, 2022; Li, 2020). At the same time, organizations are pressured to deploy software faster and more frequently through DevOps practices, while simultaneously ensuring that releases are secure, compliant, and resilient to failures (Kumar, 2019; Konneru, 2021). This confluence of speed, scale, and complexity has produced a new set of failure modes and operational demands that traditional reactive maintenance models cannot meet (Banerjee et al., 2021; Wang et al., 2022).

The literature supplied for this synthesis spans multiple, complementary domains: predictive analytics and its role in business intelligence and operations (Kumar, 2019); secure CI/CD pipeline integration (Konneru, 2021); modern data engineering with lakehouse principles (Kukreja & Zburivsky, 2021); geospatial and high-performance data handling (Li, 2020); domain-specific failure studies such as GPU failure prediction under deep learning workloads (Liu et al., 2023); applied interventions for infrastructure efficiency in sectors like water management (Kwikima et al., 2024); and broader studies on fault tolerance and redundancy strategies in cloud systems (Lin & Gupta, 2021). This heterogeneous mix affords an opportunity to craft a unified theoretical and practical approach to resilience—one that hinges on continuous telemetry, predictive modeling, secure automated pipelines, and rapid remediation orchestration.

While each domain offers mature approaches individually, the literature reveals several gaps when attempting to operationalize resilience at scale. Predictive models often remain siloed within research prototypes or offline analytics environments, disconnected from deployment pipelines and security practices (Kumar, 2019; Kukreja & Zburivsky, 2021). Similarly, security-focused CI/CD literature emphasizes scanning and gating (SAST, DAST, SCA) but less frequently addresses how security telemetry and predictive failure signals can be coalesced to inform safer rollouts and automated rollback strategies (Konneru, 2021). The GPU failure literature offers methods to predict hardware anomalies under deep learning workloads (Liu et al., 2023), but translating those insights into cross-layer operational strategies in cloud-scale environments requires careful attention to data flows, model lifecycles, and integration with orchestration systems (Lin & Gupta, 2021; Luo & Martinez, 2022).

The central problem this article addresses is the absence of a comprehensive, implementable framework that tightly integrates predictive analytics, secure CI/CD, and high-performance fault diagnosis to guide practitioners and researchers. Specifically, the literature gap is threefold: (1) lack of unified telemetry and data-curation practice tailored for model-driven operations; (2) insufficient integration between DevSecOps controls and predictive reliability signals; and (3) a need for fault-diagnosis strategies aligned with the performance characteristics and failure modes of modern GPUs and cloud storage/compute subsystems (Liu et al., 2023; Konneru, 2021; Kukreja & Zburivsky, 2021).

This article proposes a conceptual and operational framework to address these gaps. Drawing on lakehouse architectures for data engineering, DevSecOps practices for secure continuous delivery, and a layered predictive-fault-diagnosis approach for high-performance hardware, the framework delineates how organizations can design systems for anticipatory operations. The subsequent sections elaborate the methodology of synthesis, present descriptive results from the literature mapping, and offer a deep discussion of implications, limitations, and future research directions. Each claim is grounded in the provided references to ensure traceability and adherence to the supplied corpus.

METHODOLOGY

The methods used in this work are integrative and synthetic rather than experimental. The aim is to construct a rigorous, theoretically informed framework by systematically combining concepts, techniques, and empirical findings reported in the provided literature. The methodology comprises four interlinked steps: corpus characterization, thematic extraction, model-to-pipeline mapping, and synthesis of operational patterns.

Corpus Characterization. The body of literature supplied includes studies on predictive analytics for business intelligence and DevOps efficiency (Kumar, 2019), practical guides to data engineering and lakehouse architectures (Kukreja & Zburivsky, 2021), DevSecOps integration techniques for CI/CD (Konneru, 2021), geospatial high-performance computing considerations (Li, 2020), domain-specific empirical studies such as GPU failure prediction (Liu et al., 2023) and factory-grade diagnostic automation (Lulla et al., 2025), and applied empirical work in infrastructure optimization (Kwikima et al., 2024). Ancillary references cover cloud fault tolerance, AI-based failure management, and supervised/unsupervised learning techniques for anomaly detection (Zhang, 2022; Chen et al., 2021; Peterson et al., 2022; Xie et al., 2021). This heterogeneous corpus allows for cross-domain insights.

Thematic Extraction. Each reference was reviewed to extract core themes, methodologies, and key empirical or theoretical findings. From this extraction, recurring motifs emerged: the necessity of high-quality curated telemetry for model performance (Kukreja & Zburivsky, 2021; Li, 2020); the role of predictive analytics in operational decision-making (Kumar, 2019); the imperative to embed security controls within CI/CD without sacrificing deployment velocity (Konneru, 2021); and specific techniques for hardware failure prediction, especially for GPUs under intensive workloads (Liu et al., 2023). Thematic grouping enabled identification of interfaces where integration is both necessary and technically feasible.

Model-to-Pipeline Mapping. For operationalization, predictive models must be integrated into pipelines that support data ingestion, feature computation, model training, validation, deployment, and monitoring. The lakehouse paradigm provides a coherent structure to host raw telemetry and curated feature tables, addressing latency and governance concerns (Kukreja & Zburivsky, 2021). DevSecOps practices provide the controls (SAST, DAST, SCA) to ensure secure artifacts and runtime defenses (Konneru, 2021). The model-to-pipeline mapping step translates model lifecycle stages into CI/CD-friendly milestones and security gates. Particular attention was paid to the needs of GPU failure prediction: high-frequency telemetry capture, synchronization across control planes, and storage-efficient representations of time-series metrics (Liu et al., 2023; Li, 2020).

Synthesis of Operational Patterns. The final step synthesizes patterns across the corpus into an integrative framework. Patterns include layered diagnostics (combining supervised fail-predictions with unsupervised anomaly detection), early-warning scorecards for operations, security-aware deployment strategies that incorporate predictive confidence measures, and governance structures for telemetry and model validation. For each pattern, we identify prerequisites, technical trade-offs, and recommended technologies or algorithms grounded in the cited literature (Kukreja & Zburivsky, 2021; Konneru, 2021; Peterson et al., 2022).

Throughout the methodology, we adhere to two constraints: (1) the framework must be implementable using the architectures and tools discussed in the literature (e.g., lakehouse patterns, SAST/DAST/SCA tooling, supervised and unsupervised learning methods), and (2) every major inference or recommendation must be traceable to one or more references in the corpus (Kumar, 2019; Konneru, 2021; Liu et al., 2023). The result of this methodological approach is a comprehensive, literature-grounded framework that addresses telemetry, modeling, deployment, security, and remediation orchestration.

RESULTS

The literature synthesis yields a set of descriptive findings organized under five major themes: telemetry and data engineering, predictive analytics integration with operations, DevSecOps-enabled secure deployment, GPU and cloud fault diagnosis strategies, and cross-domain benefits and analogues. Each theme is presented with detailed analysis grounded in the cited works.

Telemetry and Data Engineering. A recurring and foundational finding is that predictive accuracy and operational usefulness of models are tightly coupled to the quality, fidelity, and accessibility of telemetry (Kukreja & Zburivsky, 2021). Lakehouse architectures—combining aspects of data lakes and data warehouses—offer a pragmatic path to store raw telemetry while enabling low-latency, ACID-compliant feature tables for model consumption (Kukreja & Zburivsky, 2021). For high-performance contexts (e.g., GPU clusters), telemetry must capture fine-grained time-series metrics at frequencies aligned with the dynamics of hardware degradation and workload bursts (Liu et al., 2023). Geospatial and other domain-specific high-volume data handling techniques inform scalable ingestion patterns: partitioning strategies, compression-aware formats, and distributed processing tuned for locality reduce ingestion latency and storage overhead (Li, 2020). A direct implication is that organizations must invest in data engineering practices that prioritize schema-on-read for raw traces and curated, query-optimized feature tables for analytics (Kukreja & Zburivsky, 2021; Li, 2020).

Predictive Analytics Integration with Operations. The idea of integrating predictive analytics into the operational lifecycle is strongly emphasized in the literature (Kumar, 2019). Predictive models, when embedded into decision loops, can shorten mean time to detection and mean time to repair by providing early warnings and suggested remediation actions (Peterson et al., 2022). Two complementary modeling paradigms emerge: supervised learning for known, labeled failure patterns; and unsupervised or self-supervised methods for anomaly detection that can identify previously unseen failure modes (Peterson et al., 2022; Xie et al., 2021). The pipeline must enforce continuous model validation and drift detection, because telemetry distributions shift with software and hardware upgrades, making static models brittle (Kumar, 2019). Operational integration also requires a runbook mapping: when a model emits a high-confidence failure prediction, predefined automated or semi-automated actions are invoked, balancing the cost of false positives against the risk of costly failures (Lin & Gupta, 2021).

DevSecOps-Enabled Secure Deployment. Security within CI/CD is not merely a bolt-on scanning step; rather, it must be embedded in the lifecycle to ensure that predictive models and their deployment artifacts do not introduce vulnerabilities (Konneru, 2021). SAST (static analysis), DAST (dynamic analysis), and SCA (software composition analysis) provide layered controls for code quality, runtime behavior, and third-party dependency risks (Konneru, 2021). Additionally, security telemetry—logs, scan results, and vulnerability severity metrics—should be merged with operational and failure telemetry in the lakehouse to enable holistic risk assessments. This combined view permits policies such as gating deployments when critical vulnerabilities co-occur with elevated failure risk, thus preventing exacerbation of systemic fragility (Konneru, 2021).

GPU and Cloud Fault Diagnosis Strategies. GPUs, due to their specialized memory hierarchies, thermals, and workload-dependent stress patterns, exhibit unique failure signatures (Liu et al., 2023). The literature demonstrates that deep learning workloads can precipitate hardware failures via long execution traces, memory fragmentation, and overheating, which can be detected through high-precision predictive models trained on workload-specific telemetry (Liu et al., 2023). Factory-grade diagnostic automation and domain-specific heuristics further enhance detection fidelity by incorporating manufacturing test baselines and longitudinal device histories (Lulla et al., 2025). Cloud-wide redundancy strategies and AI-optimized replication reduce data loss probability and improve tolerance to node and device failures when combined with predictive signals that prioritize preemptive migration or scaled replication (Yamamoto & Kim, 2021; Lin & Gupta, 2021). The key operational finding is that layered diagnostics—where rapid, low-cost anomaly detectors triage signals for higher-fidelity predictive models—offer a cost-effective path to operationalize GPU fault prediction (Peterson et al., 2022; Liu et al., 2023).

Cross-Domain Benefits and Analogues. Application of these techniques beyond pure compute infrastructure yields measurable benefits. For instance, integrated data-driven approaches reduced non-revenue water in water distribution networks by enabling targeted interventions derived from predictive models and curated telemetry (Kwikima et al., 2024). This cross-domain success reinforces the generality of the proposed framework: wherever reliable telemetry and secure, automated pipelines exist, predictive approaches can enable anticipatory and optimized operations (Kumar, 2019; Kwikima et al., 2024). Furthermore, high-performance geospatial data handling techniques inform the handling of large telemetry volumes, particularly in distributed systems where locality-sensitive processing reduces network bottlenecks (Li, 2020).

Synthesis: An Integrated Operational Framework. Combining the themes above suggests a coherent architecture: a telemetry-first lakehouse foundation; a continuous modeling lifecycle that includes supervised and unsupervised approaches and model governance; DevSecOps gates that ensure security and compliance of models and pipelines; and a layered diagnostic and remediation orchestration that maps predictive confidence to automated actions and human-in-the-loop escalations. This architecture reduces operational fragility by enabling early detection, secure deployments, and preemptive mitigation strategies, particularly in GPU-intensive and cloud-scale environments (Kukreja & Zburivsky, 2021; Konneru, 2021; Liu et al., 2023).

DISCUSSION

The preceding synthesis yields several deep implications, counterarguments, technical trade-offs, and areas where further theoretical and empirical work is essential. This discussion examines these aspects in depth, offering guidance for both researchers and practitioners.

The Promise of Telemetry-First Lakehouses. Lakehouse architectures reconcile the traditionally opposing goals of raw telemetry retention and performant analytics (Kukreja & Zburivsky, 2021). By preserving lineage and enabling ACID operations on curated tables, lakehouses support reproducible model training and governance—prerequisites for deploying predictive models in production. However, certain trade-offs arise: maintaining detailed telemetry at high frequencies for GPU clusters can impose storage and ingestion costs and complicate retrieval latency under peak loads (Li, 2020; Liu et al., 2023). Effective mitigation requires careful retention policies, tiered storage strategies, and compression-friendly telemetry formats. The literature suggests adjustable sampling strategies—adaptive to workload phases—to balance fidelity with cost (Li, 2020). Researchers can explore principled sampling schemes that preserve predictive features while reducing overhead.

Embedding Predictive Analytics into Operations: Opportunities and Risks. Predictive analytics deliver operational value when models are accurate, timely, and properly trusted by operational teams (Kumar, 2019). One practical challenge is model explainability: operations teams must understand why a model issued a prediction to decide whether to trust automated remediation or to intervene manually. Techniques for explainability—feature attribution, counterfactuals, and model-agnostic interpretation—should be integrated into alert payloads and runbooks to support trust (Peterson et al., 2022). Nonetheless, explainability techniques can be misleading when models face distributional shifts; claims of causality derived from observational telemetry must be handled conservatively (Kumar, 2019). Future work should investigate human-centered interfaces that combine predictive scores with confidence intervals and context-aware explanations to facilitate appropriate human–machine collaboration.

Security-Performance Trade-Offs in DevSecOps. Integrating SAST/DAST/SCA into pipelines is central to preventing vulnerabilities from reaching production (Konneru, 2021). However, security scans can create latency in rapid deployment cycles if not optimized. The recommended approach is risk-based gating: prioritize critical vulnerabilities and use staged scanning for lower-risk artifacts (Konneru, 2021). Furthermore, the literature reveals an opportunity to use predictive reliability signals to inform security decisions—for example, delaying deployment if a release coincides with elevated hardware failure probability in target clusters. Conversely, the presence of a security vulnerability in an artifact may change the operations team's willingness to perform preemptive migration, creating complex decision boundaries that require well-defined policies and decision-support tools (Konneru, 2021; Lin & Gupta, 2021).

GPU-Specific Failure Modes and Operational Strategies. GPUs differ from general-purpose servers in several respects: memory error profiles, thermal stress behavior, and sensitivity to long epochs of computation (Liu et al., 2023). Predictive models tailored to GPUs must therefore incorporate workload characteristics (e.g., batch sizes, kernel types), hardware telemetry (e.g., memory ECC events, temperature, clock throttling), and historical manufacturing/test data (Liu et al., 2023; Lulla et al., 2025). However, a potential counterargument is the risk of false positives—preemptively evacuating workloads from GPUs based on imperfect models can reduce utilization and revenue (Lin & Gupta, 2021). The layered diagnostics approach mitigates this by using conservative thresholds for automated preemption while escalating ambiguous signals for human review. Future empirical work should quantify the cost trade-offs between false positives (unnecessary migrations) and false negatives (unscheduled failures) across different workload classes.

Model Governance, Drift Detection, and Lifecycles. Models in operational contexts degrade as system characteristics evolve (Kumar, 2019). The literature underscores the need for continuous evaluation, retraining triggers based on drift detection, and rollback mechanisms in deployment orchestration (Kukreja & Zburivsky, 2021). Drift detection methods vary—from distributional divergence metrics to monitoring actionable performance metrics like recall at specific operating points (Peterson et al., 2022). A significant limitation is the lack of standardized benchmarks and telemetry schemas across organizations, complicating model portability. Researchers and standards bodies should collaborate to define open telemetry formats and benchmark datasets for fault prediction that allow reproducible evaluation and cross-validated model comparisons.

Analogues from Non-Compute Domains. The successful application of integrated data-driven approaches to reduce non-revenue water in peri-urban Tanzania demonstrates the cross-domain applicability of the framework (Kwikima et al., 2024). The analogue suggests that the primary ingredients—accurate telemetry, curated features, and operational integration—are generalizable. Yet, compute infrastructure is distinct in its temporal dynamics and scale: failures can cascade rapidly and affect global services. Therefore, while analogues are informative, domain-specific adaptation remains essential.

Governance and Organizational Readiness. Implementing the proposed framework requires organizational shifts: creation of cross-functional teams that combine data engineering, security, SRE, and hardware diagnostics expertise (Kukreja & Zburivsky, 2021; Konneru, 2021). Institutional inertia and siloed responsibilities present barriers. Incentive structures should be realigned to reward investments in reliability and security, not just feature velocity. Moreover, knowledge transfer from hardware vendors—particularly around manufacturing test baselines and warranty data—can materially improve predictive accuracy (Lulla et al., 2025). Policy frameworks that govern telemetry privacy and retention must also be considered, especially where operational telemetry overlaps with user data.

Limitations of the Current Synthesis. This article synthesizes the supplied literature but does not present new empirical experiments. While the findings are grounded in peer-reviewed and applied studies, the operational effectiveness of the integrated framework must be validated in controlled deployments. The literature selection, while diverse, may omit recent developments beyond the supplied references; for example, advances in federated model governance or latest cloud-provider-specific orchestration features may refine operational tactics (Zhang, 2022; Luo & Martinez, 2022). The synthesis also assumes access to comprehensive telemetry streams—an assumption that may not hold in legacy systems or highly regulated environments where telemetry is restricted.

Future Research Directions. Several promising directions arise from the synthesis:

1. Benchmarking and Open Datasets: Development of public, anonymized telemetry benchmarks for GPU failure prediction and cloud fault scenarios to enable standardized model evaluation (Liu et al., 2023).

2. Adaptive Sampling and Compression: Research into adaptive telemetry sampling algorithms that preserve predictive fidelity while minimizing storage and bandwidth costs (Li, 2020).

3. Explainability for Operations: Human-centered studies to determine the most useful forms of model explanations for operations teams, balancing depth with cognitive load (Peterson et al., 2022).

4. Policy-aware Deployment Orchestration: Automation frameworks that integrate security assessment, predictive confidence, and cost/risk models to make deployment decisions under uncertainty (Konneru, 2021; Lin & Gupta, 2021).

5. Cross-Organization Sharing Protocols: Secure, privacy-preserving protocols for sharing failure signatures or model artifacts between vendors and operators to accelerate detection of rare failure modes (Lulla et al., 2025).

Researchers should prioritize empirical validations that couple model performance metrics with operational KPIs, such as downtime reduction, mean time to repair, and cost impacts of preemptive interventions.

CONCLUSION

This article synthesizes a multidisciplinary literature corpus to propose an integrated, implementable framework that combines lakehouse data engineering, predictive analytics, and DevSecOps-enabled CI/CD to enhance resilience in modern computational infrastructures—particularly in environments dominated by GPU workloads and cloud-scale services (Kukreja & Zburivsky, 2021; Konneru, 2021; Liu et al., 2023). The central tenets of the framework are: (1) telemetry-first data architectures that reconcile raw trace retention with curated feature tables; (2) continuous modeling lifecycles that marry supervised and unsupervised approaches for fault prediction and novelty detection; (3) security-embedded deployment pipelines that consider predictive reliability signals alongside vulnerability assessments; and (4) layered diagnostic and remediation orchestration that balances automation with human oversight.

These components collectively address the literature gap in operationalizing predictive reliability while preserving security and deployment velocity (Kumar, 2019; Konneru, 2021). Although the framework is rooted in the extant literature rather than novel empirical trials, it provides a clear roadmap for practitioners to design systems capable of anticipatory operations and for researchers to formulate targeted validation studies. The effectiveness of this approach will depend on investments in scalable telemetry, standardized schemas, cross-disciplinary governance, and robust model governance practices. Ongoing research, benchmarking, and cross-industry collaboration will be critical to refine the proposed architecture and to ensure it remains adaptive to evolving workload characteristics and emerging failure modes (Liu et al., 2023; Lin & Gupta, 2021).

REFERENCES

1. Karwa, K. (2024). Navigating the job market: Tailored career advice for design students. International Journal of Emerging Business, 23(2). https://www.ashwinanokha.com/ijeb-v23-2-2024.php

2. Konneru, N. M. K. (2021). Integrating security into CI/CD pipelines: A DevSecOps approach with SAST, DAST, and SCA tools. International Journal of Science and Research Archive. Retrieved from https://ijsra.net/content/role-notification-scheduling-improving-patient

3. Kukreja, M., & Zburivsky, D. (2021). Data Engineering with Apache Spark, Delta Lake, and Lakehouse: Create scalable pipelines that ingest, curate, and aggregate complex data in a timely and secure way. Packt Publishing Ltd.

4. Kumar, A. (2019). The convergence of predictive analytics in driving business intelligence and enhancing DevOps efficiency. International Journal of Computational Engineering and Management, 6(6), 118-142. Retrieved from https://ijcem.in/wp-content/uploads/

5. Kwikima, M. M., Bennett, G., Ahmada, F. K., & Magina, A. (2024). Reducing non-revenue water in peri-urban Tanzania through an integrated data-driven approach: a pilot study in Dodoma. International Journal of Energy and Water Resources, 1-19.

6. Li, Z. (2020). Geospatial big data handling with high performance computing: Current approaches and future directions. High Performance Computing for Geospatial Applications, 53-76. AMERICAN ACADEMIC PUBLISHER. https://www.academicpublishers.org/journals/index.php/ijvsli

7. Liu, H., Li, Z., Tan, C., Yang, R., Cao, G., Liu, Z., & Guo, C. (2023, June). Predicting GPU Failures With High Precision Under Deep Learning Workloads. In Proceedings of the 16th ACM International

Conference on Systems and Storage (pp. 124-135).

8. Zhang, K. (2022). Cloud Computing in Modern IT Infrastructure. IEEE Transactions on Cloud Computing, 10(3), 456-468.

9. Chen, M., Zhang, L., Li, Y., & Hu, S. (2021). AI in Cloud Fault Tolerance: A Comprehensive Survey. Journal of Cloud Engineering, 8(2), 123-138.

10. Patel, R., & Singh, T. (2022). Failure Detection in Cloud-Based Services Using AI and Machine Learning. ACM Computing Surveys, 54(5), 1-28.

11. Banerjee, S., Kumar, A., & Lee, J. (2021). A Study on Traditional vs. AI-Based Fault Tolerance Mechanisms in Cloud Computing. Future Generation Computer Systems, 127, 89-104.

12. Wang, H., et al. (2022). Self-Healing Cloud Systems: The Role of AI and ML in Proactive Failure Management. IEEE Transactions on Dependable and Secure Computing, 19(2), 289-306.

13. Lin, J., & Gupta, P. (2021). AI-Optimized Redundancy Strategies for Cloud Computing. Journal of Parallel and Distributed Computing, 155, 150-165.

14. Luo, C., & Martinez, R. (2022). Google Cloud's AI-Based Fault Tolerance: An Empirical Analysis. IEEE Cloud Computing, 9(3), 67-79.

15. Yamamoto, T., & Kim, S. (2021). Reducing Data Loss Probability in Cloud Storage Using AI-Enhanced Replication. ACM Transactions on Storage, 17(2), 1-19.

16. Peterson, K., et al. (2022). Supervised Learning Techniques for Predictive Failure Analysis in Cloud Computing. IEEE Transactions on Ne twork and Service Management, 18(4), 512-530.

17. Xie, Y., Huang, L., & Li, G. (2021). Unsupervised Learning for Cloud Anomaly Detection: A Case Study with Autoencoders. Journal of Cloud Security, 11(3), 129-144.

18. Lulla, K., Chandra, R., & Ranjan, K. (2025). Factory-grade diagnostic automation for GeForce and data centre GPUs. International Journal of Engineering, Science and Information Technology, 5(3), 537-544.