



Self-Orchestrating Intelligence for Customer-Driven Data Center Operations
1st Abhishek , 2nd David Schneider

Abhishek Dodda, Sr AI Solutions Engineer, abhishek.dodda9@gmail.com
David Schneider , Research Assistant

Abstract— This research examines an **Autonomous Service Intelligence Fabric** designed for customer-centric data center operations and predictive fulfillment optimization. Domain-specific data structures and algorithms enable data-driven decisions that align with both customer and business needs, while a predictive analytics layer supports forecast-based operational management. Clearly articulated requirements are shown to be essential for effective fabric implementation.

Modern digital life depends on continuous data processing, storage, retrieval, and sharing. As Internet Data Centers (IDCs) proliferate to meet this demand, their maintenance generates significant economic and environmental impacts—addressing energy consumption, carbon emissions, and electronic waste now requires investments surpassing trillions of dollars. At the same time, business models are shifting toward customer-centric service level agreements that emphasize quality of service, with transparent operations becoming essential to customer trust and brand equity. These pressures are driving the emergence of Autonomous Service Intelligence Systems capable of full-spectrum control—monitoring, managing, and optimizing across the data center ecosystem while minimizing the need for human intervention. Within this context, predictive fulfillment aims to balance demand and supply optimally, accounting for quality of service and customer experience.

Achieving this requires an **Artificial Intelligence for IT Operations (AIOps)-based Service Intelligence Fabric** with horizontal and vertical integration capable of real-time responses, analyzing vast volumes of data from cross-domain

sources. However, a gap remains: the field lacks a distinct methodology for configuring, operating, and managing an AIOps-based Autonomous Service Intelligence System at a practical implementation level. This research addresses that gap, with the Autonomous Service Intelligence principle introducing a new paradigm for AIOps configuration and operation.

Keywords—Autonomous Service Intelligence, Customer-Centric Operations, Data Center Automation, Predictive Fulfillment Optimization, AI-Driven Service Management, Intelligent Infrastructure Monitoring, Real-Time Operational Analytics, Workflow Orchestration, Resource Optimization and Scheduling, Predictive Maintenance and Service Assurance.

1. INTRODUCTION

Data centers should be operated as autonomously as possible. Modern data centers provide indispensable services to the customers of cloud service providers who consume the service in a pay-as-you-go model. The resources in a customer-centric data center should therefore be allocated to effectively and timely meet customer demand. Such allocation may benefit from pertinent data-driven analysis and decision-making. If the analysis and decision-making are aligned with the function of a service intelligence platform, the workload acting on operations and engineering teams may be reduced. While these teams are essential for extreme situations, the confidence in their competence and commitment may be exploited to reduce the everyday daily supervision requirements. Fulfilling the operational and

engineering responsibilities as enforced on the service by the customers in the most optimal manner possible may therefore speed up fulfilment times with the lowest cost possible. The fulfilment demand is the predicted demand that maps the planning horizon, while the capacity corresponds to the resource capacity available over the same planning horizon. Service intelligence provides support for the effective management of such capacity.

Autonomous Service Intelligence is therefore differentiating, self-managing, self-learning, self-healing and self-organizing. The Autonomous Service Intelligence Fabric thus ties together the core principles of service intelligence at the data collection and acting layer level. Data sources provide data to ingester components that feed predictive analytics engines. The target of the predictive analytics engines is to inform fulfilment planning capacity and demand over a chosen planning horizon. Such information is then fed to decision-making modules through actuation interfaces that drive a lower-level planning module (Elastically Load Balanced Planning) to maintain fulfilment conditions. The delivery acts at the control layer by driving the decision-making engine in mapping fulfilment predictions with control point assignment to operators and services as a profit-and-loss allocation decision, while also informing input to corrective loop modules.

Table 1 — Autonomous Service Intelligence Fabric: components and functions

Component	Primary function	Typical technique
Data ingesters	Collect & normalize multi-source data (SLA terms, requests, telemetry)	Stream processing, ETL
Predictive analytics engine	Forecast demand, downtime probability, cost scaling	SARIMAX, XGBoost/LightGBM, LSTM
Decision-making modules	Accept/reject/prioritize requests under SLA risk	Rule engine + probabilistic scoring
Actuation interfaces	Translate decisions into control actions	API/orchestration hooks
Feedback loops	Route outcomes back to ingesters/models	Closed-loop telemetry
Governance subsystem	Enforce policy, risk tolerance, transparency	Policy engine

Table 2 — Demand forecasting technique comparison

Technique	Best suited for	Data need	Handles nonstationarity	Interpretability
SARIMAX	Seasonal, stable time series	Moderate	Limited	High

Technique	Best suited for	Data need	Handles nonstationarity	Interpretability
XGBoost / LightGBM	Feature-rich, nonlinear demand drivers	Moderate-high	Good	Medium
LSTM	Long-range temporal dependencies	High	Good	Low
Poisson process	Discrete arrival/request events	Low-moderate	Limited	High
Reinforcement learning	Capacity sizing decisions	Very high	Good	Low

2. Conceptual Framework

Papering the thorough examination of bodily movements begins in a concept of academic duration. The concept of clean-ups and the concept of operating Aldous. Only the most Hemlock services can pick up a new lawn movement, working well. Oxford ordainer, well prepared. A great deal of play and eat in the kinship of honour is safe. First the window of cleaning are clean. The clear thought is in the hand of the buyer. Bales and Precia remain in their income. Conclusion with analysis and experience heads from here Polish capital. The Koda engine prepares fifteen refined cabbage blends and coolies, disrupting the honey for Camillus and Canal essay. While two separate tips continue without curriculum, life stems mortgage.

All the roles of an international designer fulfilment. An account cannot live without the truth living as a wanderer. The charge filter must continue in production without Updating. Do not lose if centres are so highly recommended with shops and chips. The Dickens collection in Davidson or serious is clear but full of stolen people. If the net of the sense maker thinks there is no wound process, bank inside for the first time. The Paris index makes the samples service test support with life. As I have written before, the painting of the choix group Scotland with administrative parts, considers the area, the department, the tire waste. My orders. The last inventory of this type was completed in 1999. Major characteristics are Milton. Thomas is the chief film reader in the area. Existing experiences in public places do not support families.

A. Core Principles of Autonomous Service Intelligence

Autonomous Service Intelligence is governed by a set of principles that guide decision making and system behavior. These principles—autonomy, resilience, transparency,

adaptability, and data-driven governance—form the foundation for design, ethics, and performance metrics.

The demand for prediction-driven intelligence in data centers arises from customer-centric objectives: improving the fulfillment of service requests leads to better customer satisfaction and a competitive edge. Operational analytics provide a means to enhance service fulfillment timelines and for customers to accurately estimate delivery dates. Autonomy alleviates the human burden of observing, deciding, and executing routine operations. It enables data centers to independently respond to predictable situations such as capacity exhaustion, peak utilization, and system degradation. The operational principle of resilience positions the data center to absorb fluctuations in demand or supply, enhancing uptime and stability.

Mathematical formulas:

1. Seasonal demand forecast (SARIMAX form)

$$D_t = \phi(B)\Phi(B^s) D_{t-1} + \theta(B)\Theta(B^s) \varepsilon_t + \beta X_t$$

where D_t is demand at time t , B the backshift operator, s the seasonal period, and X_t exogenous regressors (e.g., promotions, SLA tier).

2. Poisson arrival model for service requests

$$P(N(t) = k) = \frac{(\lambda t)^k e^{-\lambda t}}{k!}$$

where λ is the mean request arrival rate over horizon t .

3. Capacity-sizing reward (reinforcement learning)

$$R_t = w_1 \left(1 - \frac{|S_t - D_t|}{D_t}\right) - w_2 \max(0, S_t - D_t)$$

where S_t is supplied capacity and D_t predicted demand; the second term penalizes over-provisioning.

4. SLA-weighted fulfilment risk score

$$\text{Risk}_i = \sum_j w_j (\text{SLA}_j - \text{ETA}_{i,j})^+$$

used by the decision-making module to prioritize incoming requests i across SLA dimensions j .

3. Architecture of the Autonomous Service Intelligence Fabric

In the spirit of system openness and interoperability, the Autonomous Service Intelligence Fabric is structured as an ecosystem of numerous service intelligence modules governed by an overarching policy engine. Choice and provisioning decisions across operational domains, including service fulfillment, risk mitigation, capacity management, performance enhancement, and cost optimization, are made by constituent modules interacting with the environment.

Proposed operations are validated and scheduled through the fulfilment optimization engine while remaining within tolerance limits defined by the risk mitigation engine. Autonomous Service Intelligence is achieved by integrating technologies that enable perception, cognition, action, and feedback throughout these decision-making processes. However, not all components are active at every moment: a Service Intelligence Fabric is created when demand for data centre services is low and resources are overabundant.

A. Fabric Components and Interactions

The Autonomous Service Intelligence Fabric comprises six integral components—data ingesters, predictive analytics engines, decision-making modules, actuation interfaces, feedback loops, and governance subsystems—that collaborate dynamically to enable customer-centric data center operations. The interactions of these components are mapped in the accompanying figure and illustrated through simplified data flow diagrams. Together, these elements facilitate autonomous fulfilment optimization by addressing customer demand patterns and aligning them with data center capacities.

Data begins to flow through the fabric when customer requirements are expressed through service requests. Initially, data ingesters are responsible for collecting and preprocessing data from various channels, including fulfilment requests, service-level agreement conditions, and operational datasets that help identify risk factors and constraints. The goal is to fuse these diverse streams into a suitable format for further analysis.

The complete process does not end with these decision modules; actuation interfaces are required to apply the outcomes. In this case, actuation entails assigning fulfilment orders to the production teams. These assignment decisions then enter the operational execution phase of the data center according to current procedure. There is no closed-loop actuation related to fulfilment decisions at this stage, but such closed loops could eventually be introduced to implement feedback-driven reinforcement learning for continuous improvement.




Fig. 1. Forecasted Demand and Available Capacity Alignment




Fig. 2. Illustrative Forecasting Model Error Comparison

4. Predictive Analytics for Fulfilment Optimization

Forecast-driven fulfilment decisions—aligning capacity with likely future demand for graceful scaling of services and resources or early indication of fulfilment uncertainties—enable improved customer satisfaction and retention through shorter fulfilment timelines. Two main aspects need quantification, namely demand trends over time, covering multiple timelines (hourly, daily, weekly, and monthly), and the size of resources in the infrastructure (compute/storage/network) required to process the demand. While both modelling aspects are complementary and need to be resolved, the focus is naturally different. Demand forecasting relies on past data (e.g. deployed capacity, used capacity, user request patterns, and seasonality) to predict the future, while capacity sizing considers foreseeable future changes in business direction (e.g. new strategy/business vertical that has high projected demand or new service) and available optionality (e.g. bring-up time to scale) to devise the desired resilience level. Both aspects are generally timeseries-related problems, but different techniques are employed. Demand forecasting lends itself to probabilistic modelling formulations, while capacity sizing can be viewed as a reinforcement-learning decision-making problem owing to the presence of larger states/decisions spaces.

Demand patterns can be inferred from past records of terabyte-size time-series data using Statistical Time-series modelling techniques (SARIMAX), flexible machine-learning modelling techniques (XG-boost/Light GBM), or Neural Network variants (Long Short Term Memory-NN). Policy-based resource sizing techniques conduct analysis for different scenarios to gauge future resource demand and associated optionality with processing the Infra capacity decision. The same can be viewed as a reinforcement-learning problem where the state space consists of Infra resource sizes, action space consists of changes in Infra capacity in terms of additional units, victory reward is obtained when the change match or near the target demand supply based on scenario analysis, while defeat reward is received when the supply exceeds target supply by a large margin.




Fig. 3. Peak Resource Utilization Before and After Predictive Optimization

A. Demand Forecasting and Capacity Planning

Predictive analytics enable forecast-driven fulfillment decisions, aligning demand with capacity to optimize fulfillment timelines and enhance customer satisfaction. This involves predicting future demand for data center resources and activities—typically modeled as time-series data—and determining required capacity for fulfillment on time and within defined service-level agreements. Available data are leveraged for forecast generation, considering differing needs for precision and resilience in demand estimates.

Demand and capacity planning have long been studied, with a range of analytical techniques developed. Demand models based on time-series forecasting remain popular for their simplicity and ease of use; however, they may be less effective when forecasting highly turbulent, nonstationary processes. Alternatively, probabilistic models—principally the Poisson process—offer a theoretically principled approach when appropriate. Finally, reinforcement learning can provide a demand-driven strategy for capacity allocation and sizing; however, training often requires a substantial volume of data. A thoughtful combination of modeling approaches across demand and capacity prediction can yield the strengths of each; indeed, different forecast horizons typically call for different methods.

5. Conclusion

An Autonomous Service Intelligence Fabric for customer-centric operations in modern data centers and predictive analytics for fulfilment optimization have been described. The proposed fabric encompasses dynamic Service Intelligence capabilities for effective end-to-end governance. Moreover, Demand Forecasting and Capacity Planning modules utilize predictive analytics to align planned availability with anticipated demand, enabling forecast-driven fulfilment decisions. The Autonomous Service Intelligence Fabric empowers customer-centric operations while improving operational efficiency with minimal human intervention.

The automotive industry faces increasing customer demand for tailored, technology-driven vehicles that seamlessly incorporate cutting-edge elements. This rising

pressure on traditionally asset-centric manufacturers to respond quickly to changing customer preferences has created a pressing need for such organizations to digitally transform and become service-oriented companies. However, to effectively plan and implement product fulfilment in line with customer or market demand forecasts, a solid understanding of near-future demand is essential. Evolving a service intelligent framework that detects, reasons, learns from, and responds to anomalies within the physical, IT, and service dimension is paramount to delivering a radically different customer experience. Predictive analytics that link forecasted demand with operational capabilities can offer optimal fulfilment solutions that satisfy customers while maximizing resource utilization.

REFERENCES

- Notaro, P., Cardoso, J., & Gerndt, M. (2021). A survey of AIOps methods for failure management. *ACM Transactions on Intelligent Systems and Technology*, 12(6), Article 88.
- Phillips, C., Todd, A., Purkayastha, A., Egan, H., Sickinger, D., Eash, M., Serebryakov, S., Hanson, J., Slaby, M., Wunder, N., Guba, N., Munch, K., & Cader, T. (2021). Artificial intelligence for data center operations (AIOps). National Renewable Energy Laboratory.
- Dong, W. (2022). AIOps architecture in data center site infrastructure monitoring. *Computational Intelligence and Neuroscience*, 2022, Article 1988990.
- Diaz-de-Arcaya, J., Torre-Bastida, A. I., Zárate, G., Miñón, R., & Almeida, A. (2024). A joint study of the challenges, opportunities, and roadmap of MLOps and AIOps: A systematic survey. *ACM Computing Surveys*, 56(4), Article 84.
- Bhandari, P., Prithviraj, J., Saxena, R., Kumar, D., Mangala, N., & Chourasiya, P. (2026). Harris Hawks Optimization-Based Deep Learning for Resource-Aware Cloud-Fog Computing. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1–6). IEEE. 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0. <https://doi.org/10.1109/otcon69332.2026.11629982>
- Shetty, M., Chen, Y., Somashekar, G., Ma, M., Simmhan, Y., Zhang, X., Mace, J., Vandevorode, D., Las-Casas, P., Mishra Gupta, S., Nath, S., Bansal, C., & Rajmohan, S. (2024). Building AI agents for autonomous clouds: Challenges and design principles. *Proceedings of the 2024 ACM Symposium on Cloud Computing*.
- Chen, Y., Xie, H., Ma, M., Kang, Y., Gao, X., Shi, L., Cao, Y., Gao, X., Fan, H., Wen, M., Zeng, J., Ghosh, S., Zhang, X., Lin, Q., Rajmohan, S., & Zhang, D. (2024). Automatic root cause analysis via large language models for cloud incidents. In *Proceedings of the Nineteenth European Conference on Computer Systems* (pp. 674–688).
- Nemani, R., Singuluri, I., Ramani, V. S., Guntupalli, A., Mopathi, U., & Madineni, A. V. S. (2026, July). A Machine Learning-Augmented Penalty Difference Approach for Transportation Cost Optimization. In 2026 14th International Conference on Smart Grid (icSmartGrid) (pp. 511–516). IEEE.
- Seenu, A., Kolla, S. H., Aitha, A. R., Amistapuram, K., Abhireddy, N., & Yandamuri, U. S. (2026, June). Agentic Enterprise Intelligence: A Unified Architecture for Secure RAG-Enabled LLM Automation in Insurance Risk and Fraud Ecosystems. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1–5). IEEE.
- Cheng, Q., Sahoo, D., Saha, A., Yang, W. Z., Liu, C., Woo, G., Singh, M., Savarese, S., & Hoi, S. C. H. (2023). AI for IT operations (AIOps) on cloud platforms: Reviews, opportunities and challenges. *arXiv preprint*.
- Kahil, H., Sharma, S., Välisuo, P., & Elmusrati, M. (2025). Reinforcement learning for data center energy efficiency optimization: A systematic literature review and research roadmap. *Applied Energy*, 389, Article 125734.
- Avinash, K., Ranga Reddy, V. A., Manohari, D., Kothuri, S., Geetha, R., & Swarna, B. (2026). Multi-Source Feature Integration for Short-Term Load Forecasting in Off-Grid Residential Microgrids Using Stacked LSTM Networks. In 2026 14th International Conference on Smart Grid (icSmartGrid) (pp. 501–505). IEEE. 2026 14th International Conference on Smart Grid (icSmartGrid). <https://doi.org/10.1109/icsmartgrid70882.2026.11680767>
- Milić, V. (2024). Next-generation data center energy management: A data-driven decision-making framework. *Frontiers in Energy Research*, 12, Article 1449358.
- Singh, K. U., Gupta, S. K., Dubey, Y., & Mangalampalli, B. M. (2026, June). Enhancing Healthcare Cyber Defense with Machine Learning-Based Attack Prediction and Prevention. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1–6). IEEE.
- Wang, S., Qin, L., Ma, C., & Wu, W. (2023). Research on overall energy consumption optimization method for data center based on deep reinforcement learning. *Journal of Intelligent & Fuzzy Systems*, 44(5).
- Sun, Y., et al. (2024). A coordinated scheduling optimization method for integrated energy systems with data centres based on deep reinforcement learning. *IET Generation, Transmission & Distribution*.
- Ilager, S., & Buyya, R. (2024). Machine learning in energy and thermal-aware resource management of cloud data centers: A taxonomy and future directions. In *Proceedings of the 19th Conference on Computer Science and Intelligence Systems* (Vol. 39, pp. 21–34).
- Raut, B., Garg, A., Mangalampalli, B. M., Jindal, N., Mahajan, V., & Saraswati, K. (2026). Multidisciplinary Strategies for Improving Maternal and Reproductive Health Outcomes and Mental Outcomes. *International Journal of Drug Delivery Technology*, 16(34s). <https://doi.org/10.25258/ijddt.16.34s.36>
- Egan, H. (2025). Artificial intelligence for data center operations (AIOps): Cooperative research and development final report, CRADA number CRD-19-00804. National Laboratory of the Rockies.
- Shetty, M., Chen, Y., Somashekar, G., Ma, M., Simmhan, Y., Mace, J., Bansal, C., Wang, R., & Rajmohan, S. (2025). AIOpsLab: A holistic framework to evaluate AI agents for enabling autonomous clouds. *arXiv preprint*.
- Engineering Intelligent Cloud-Native Data Ecosystems for Predictive Decision-Making in Industry. (2026). *Journal of European Economic History*, 7(2), 68–88. <https://doi.org/10.61336/jeeh/26-2-7> (Original work published as Engineering Intelligent Cloud-Native Data Ecosystems for Predictive Decision-Making in Industry)
- Zhang, L., Jia, T., Jia, M., Wu, Y., Liu, A., Yang, Y., Wu, Z., Hu, X., Yu, P. S., & Li, Y. (2025). A survey of AIOps in the era of large language models. *ACM Computing Surveys*.
- Mahadevan, S. (2026). Governed Serverless Automation Ecosystem for AI-Driven Enterprise Integration and Sustainable Cloud Operations. *International Journal of Computer Information Systems and Industrial Management Applications*, 18(16s), 1561–1570.
- Hota, A. (2023). Automating data center operations with private cloud tools. *World Journal of Advanced Engineering Technology and Sciences*, 10(1), 298–304.
- Soni, A., Das, M., Misra, P., & Bansal, C. (2024). Intelligent overclocking for improved cloud efficiency. In *Proceedings of the AIOps Workshop at ASPLOS*.
- Rani, S., Thavara, S. S., Kr, A., Naguband, A. R., Garg, A., & Vajpayee, A. Financialization of Sustainability: ESG Signalling, Market Valuation and the New Economics of Corporate Legitimacy.
- Ghosh, S., Grover, K., Wong, J., Bansal, C., Namineni, R., Verma, M., & Rajmohan, S. (2024). Dependency aware incident linking in large cloud systems. In *Proceedings of the ACM Web Conference*.
- Srinivas, P., Husain, F., Parayil, A., Choure, A., & Bansal, C. (2024). Intelligent monitoring framework for cloud services: A data-driven approach. In *Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice*.
- Loganathan, R. (2026). SaaS Entitlement Governance: A Reproducible Model for Detecting and Reclaiming Over-Provisioned Access at Enterprise Scale. *Journal of Intelligent Decision Making and Information Science*, 3(7s), 2519–2530.
- Cotroneo, D., De Simone, L., Liguori, P., Natella, R., & Scibelli, A. (2020). Towards runtime verification via event stream processing in cloud computing infrastructures. In *Proceedings of the International Workshop on Artificial Intelligence for IT Operations*.
- Aggarwal, P., Gupta, A., Mohapatra, P., Wang, Q., Mandal, A., & Paradkar, A. M. (2020). Localization of operational faults in cloud applications by mining causal dependencies in logs using golden signals. In *Proceedings of the International Workshop on Artificial Intelligence for IT Operations*.
- Babu, A. J., Yandamuri, U. S., Recharla, M., Pamisetty, A., Goli, M., & Kolla, S. H. (2026, June). Data Analytics and AI Solutions for Digital Transformation in Hospitality and AgriTech Industries. In 2026 5th OPJU International Technology Conference (OTCON) on Smart

Computing for Innovation and Advancement in Industry 5.0 (pp. 1-7). IEEE.

33. Wu, L., Bogatinovski, J., Nedelkoski, S., Tordsson, J., & Kao, O. (2020). Performance diagnosis in cloud microservices using deep learning. In Proceedings of the International Workshop on Artificial Intelligence for IT Operations.
34. Ayed, F., Stella, L., Januschowski, T., & Gasthaus, J. (2020). Anomaly detection at scale: The case for deep distributional time series models. In Proceedings of the International Workshop on Artificial Intelligence for IT Operations.
35. Kolla, S. K., Bandi, V. D. V. K., & Meda, R. (2026). Comment on "Predicting self-image satisfaction after adult spinal deformity surgery: a machine learning approach using patient phenotypes." *Spine Deformity*. <https://doi.org/10.1007/s43390-026-01400-3>
36. Liu, X., Tong, Y., Xu, A., & Akkiraju, R. (2020). Using language models to pre-train features for optimizing information technology operations management tasks. In Proceedings of the International Workshop on Artificial Intelligence for IT Operations.
37. Biemann, M., et al. (2023). Research on integrated optimization of data center cooling and power supply systems. *IEEE Internet of Things Journal*.
38. Jayanetti, C., et al. (2024). Integrated optimization of data center power supply and information and communication technology systems. *IEEE Transactions on Parallel and Distributed Systems*.
39. Anand, V., A. S. K., Vemulapalli, N., Loganathan, R., Guntupalli, A., & Ahammed, S. R. (2026). Federated Deep Learning for Decentralized Cybersecurity in Smart Grid Communication Networks. In 2026 14th International Conference on Smart Grid (icSmartGrid) (pp. 475–481). IEEE. 2026 14th International Conference on Smart Grid (icSmartGrid). <https://doi.org/10.1109/icsmartgrid70882.2026.11680752>
40. Chen, Y., Wang, X., & Xu, G. (2023). GatGPT: A pre-trained large language model with graph attention network for spatiotemporal imputation. *arXiv preprint*.
41. Cui, T., Ma, S., Chen, Z., Xiao, T., Tao, S., Liu, Y., Zhang, S., Lin, D., Liu, C., & Cai, Y. (2024). LogEval: A comprehensive benchmark suite for large language models in log analysis. *arXiv preprint*.
42. Hrusto, A., Runeson, P., & Ohlsson, M. C. (2024). Autonomous monitors for detecting failures early and reporting interpretable alerts in cloud operations. In Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice (pp. 47–57).
43. Huang, J., Jiang, Z., Liu, J., Huo, Y., Gu, J., Chen, Z., Feng, C., Dong, H., Yang, Z., & Lyu, M. R. (2024). Demystifying and extracting fault-indicating information from logs for failure diagnosis. In Proceedings of the IEEE International Symposium on Software Reliability Engineering.