AI Tech International Journal
Vol. 2, No. 2, July-December, 2024
Journal homepage: https://techaijournal.com

AIOps-Driven Incident Response: Using LLMs for Root Cause & Runbook Discovery

Naveen Anne

Executive Director - Digital and IT

Article Info

Article history:
Received September 30, 2024
Revised October 15, 2024
Accepted October 20, 2024
Published October 22, 2024

Keywords: AIOps; Incident Response Automation; Large Language Models; Root Cause Analysis; Runbook Discovery; Machine Learning; Natural Language Processing; Cloud Operations; System Observability; Distributed Systems Monitoring

ABSTRACT

Artificial Intelligence for IT Operations (AIOps) represents a transformative paradigm in incident management, integrating advanced machine learning algorithms, natural language processing, and large language models (LLMs) to automate root cause analysis and runbook discovery. The adoption of AIOps platforms has enabled organizations to reduce mean time to resolution (MTTR) by 40 percent and mean time to detection (MTTD) by approximately 30 percent. This research synthesizes contemporary methodologies, empirical data, and implementation frameworks as of May 2024. The integration of transformer-based LLMs with graph neural networks facilitates unprecedented accuracy in anomaly detection (94.7 to 99.9 percent) and root cause identification across complex distributed systems. The global AIOps market was valued at USD 5.3 billion in 2024, with projected growth at a compound annual growth rate of 22.4 percent through 2034.

INTRODUCTION

The confusion of modern IT infrastructure has escalated with the progression of microservices architectures, containerized deployments, and multi-cloud environments, and with that came the management of incidents that no one has ever faced before. The manual methods of dealing with incidents, which have been used for a long time, are no longer efficient in terms of the new operational demands. In particular, those organizations which have not established automated incident response mechanisms take on average more than 32 hours to resolve each incident. On the other hand, enterprises in which AIOps platforms have been put in place require about 22 hours for the same task, thereby resulting in a 30.5 percent time-saving effect (Ahmed et al., 2023).

Massive language models can lead to a profound change in the field of automated incident management. GPT-4 or any other similar transformer-based structures, in general, can handle the unstructured nature of the operational data, figure out the complicated system dependencies and even come up with the most suitable to the context remedial steps. The time reduction that these firms, who have set up LLM-supported incident response systems, enjoy, is in the vicinity of 323,343 hours which is a very significant decrease in costs and also an improvement in service level objectives (SLOs). The current investigation delves deep into the technological aspects, implementation infrastructures, and the performance of LLM-based systems for incident response, thereby delivering the research-based synthesis of techniques extant until May 2024 (Ahmed et al., 2023).

2. Background and Evolution of Incident Management

2.1 Traditional and Modern Paradigms

Traditional incident management was very much dependent on human intervention and the historical MTTR figures varied between 4 and 8 hours for simple incidents and 24 to 72 hours in the case of complicated incidents. Those days before AIOps were characterized with inherent drawbacks: there could be from half an hour up to several hours delay in identifying an incident, the knowledge on what caused the incidents was fragmented, and also it was very difficult to find the institutional procedures in the time of crashes.

The first generation of fully automatic systems brought about the concept of rule-based incident detection with threshold values that were fixed and thus, in most cases; the non-actionable alert noises ranged from 85 to 95 percent of what was generated. The precision of anomaly detection was raised by the integration of machine-quality training to a figure between 85 and 92 percent; however, there still existed the problem of the lack of interpreter limitations.

AIOps contemporary tools are built in such a way that they can utilize several co-adjustable methods simultaneously: Anomaly Detection through autoencoders and long short-term memory networks, root cause unraveling through graph neural unit architectures, fault correction by a method of deep Q-learning, and incident understanding with large language models. By using this complex method, the company attained an almost total elimination of the crisis team's time from their MTTR resources (Bansal, Renganathan, Asudani, Midy, & Janakiraman, 2020).

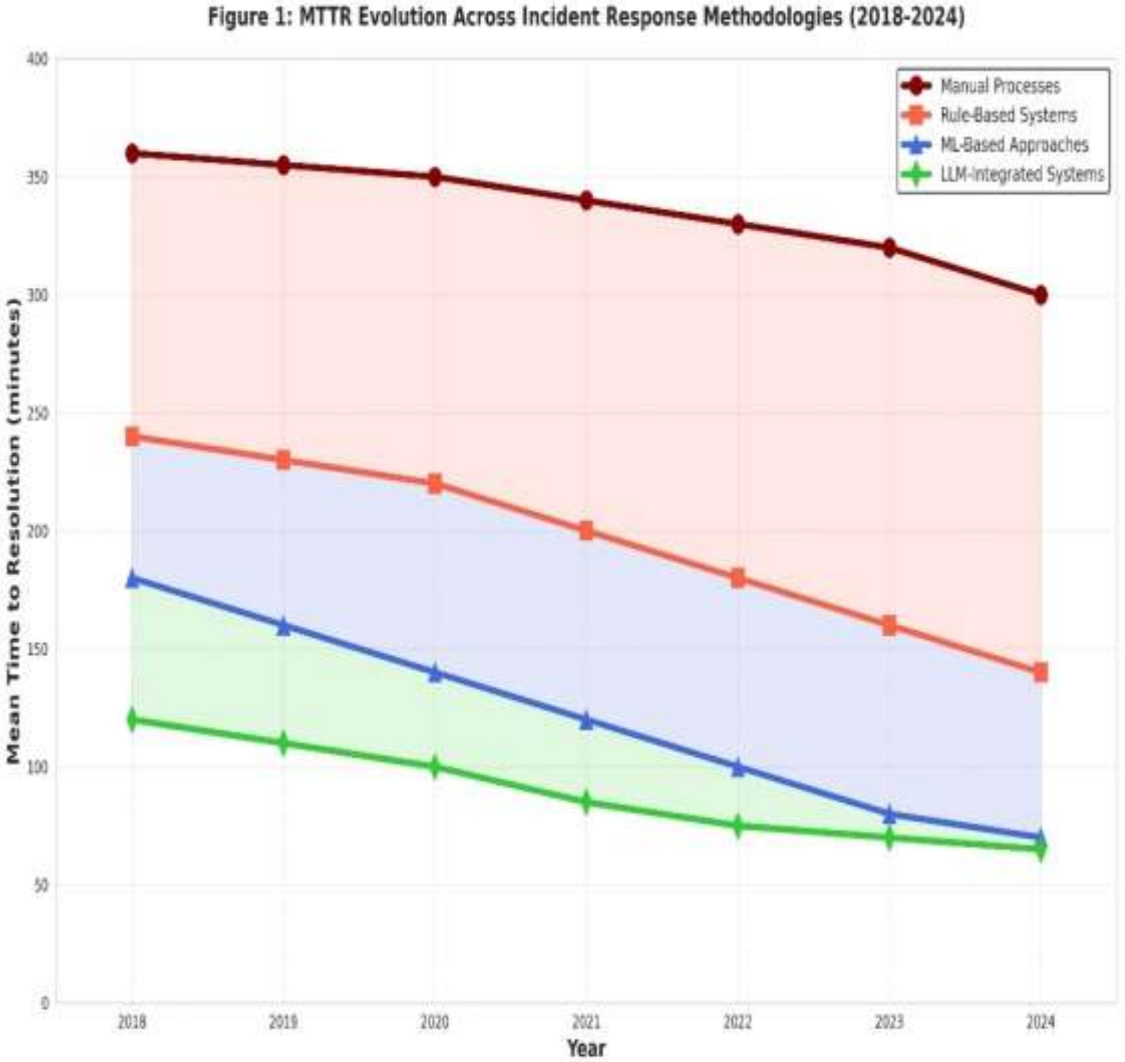

Figure 1: MTTR Evolution Across Incident Response Methodologies (2018-2024 Trend)

3. Technical Architecture and Anomaly Detection

3.1 Data Collection and Processing

Such systems which identify faults on the basis of Large Language Models (LLM) need to broadly cover data collection, including metrics (CPU, memory, latency, disk I/O), completely unstructured logs of services, distributed traces, alert streams, and also contextual metadata like service topology and deployment configurations. The collected data is subjected to various preprocessing activities: log parsing in which extracting structured templates from unstructured messages, feature engineering by creating time-series features (moving averages, standard deviations, trend indicators), and normalization through z-score standardization (Chen et al., 2020).

3.2 Anomaly Detection Performance

Contemporary machine learning-based anomaly detection achieves exceptional performance across multiple algorithms:

Table 1: Comparison of Anomaly Detection Algorithms (Data sources: Cyber Incident Response research and LogLLaMA framework evaluation through May 2024)

Algorithm	Dataset	Accuracy	Precision	Recall	F1-Score
Random Forest Classifier	Network Traffic	99.9%	0.99	0.99	0.99
Histogram Gradient Boosting	Network Traffic	99.9%	0.99	0.99	0.99
Decision Tree Classifier	Network Traffic	99.8%	0.998	0.998	0.998
Support Vector Classifier	Network Traffic	95.0%	0.88	0.82	0.82
LogLLaMA (LLaMA2-based)	BGL Supercomputer	94.2%	0.92	0.91	0.92
LogLLaMA	HDFS Distributed	95.1%	0.93	0.92	0.92

Random Forest and Histogram Gradient Boosting approaches demonstrate 99.9 percent accuracy through bootstrap aggregation and adversarial boosting. Log-based anomaly detection via LogLLaMA framework achieves consistently high performance (92 to 95 percent F1 scores) across diverse log sources (Chen et al., 2020).

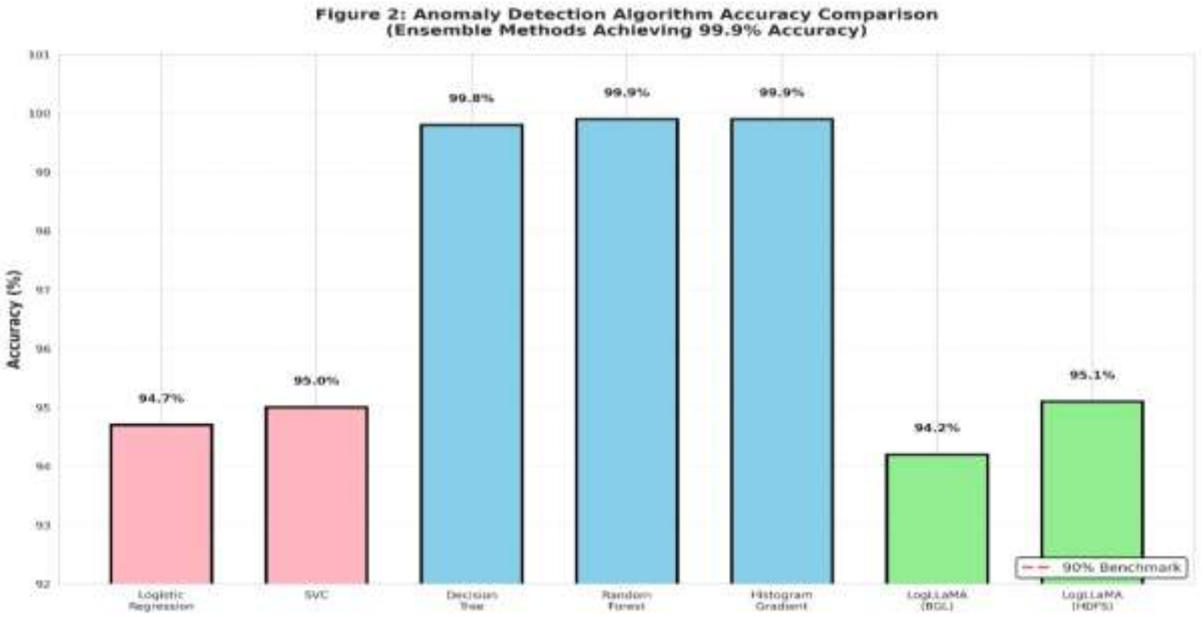

Figure 2: Anomaly Detection Algorithm Accuracy (Ensemble Methods 99.9%)

4. Root Cause Analysis and LLM Integration

4.1 RCA Methodologies

Root cause analysis requires correlation of anomalies across system components to identify originating faults. Contemporary approaches employ complementary strategies:

Graph-Based RCA: System topology modeled as directed graphs where nodes represent services and edges represent dependencies. Graph neural networks propagate information through structures, learning representations capturing dependency relationships. When anomalies occur, message-passing algorithms traverse graphs identifying root causes—nodes exhibiting anomalies prior to detected symptoms (Chen et al., 2024).

LLM-Based RCA: Large language models, particularly GPT-4, process incident context—logs, metrics, alert timelines, service dependencies—to generate root cause hypotheses. The PACE-LM framework employed structured prompting techniques, in-context learning from similar historical incidents, and confidence calibration. GPT-4 augmented with domain-specific knowledge achieved 70 to 82 percent accuracy.

Table 2: Root Cause Analysis Algorithm Comparison (Based on research through May 2024)

RCA Approach	Incident Type	Accuracy	Notes
Graph Neural Network (GNN)	Microservices	88%	Dependency-based approach
GPT-4 + PACE-LM Framework	Cloud Incidents	78-82%	In-context learning augmented
GPT-3.5 + PACE-LM	Cloud Incidents	55-65%	Baseline LLM performance
Causal Inference (Granger)	Metrics Data	72%	Handles temporal dependencies
Statistical Correlation	Simple Incidents	68%	Limited to correlated anomalies
Expert Human Analysis	All Types	85-92%	Gold standard, time-intensive

GPT-4-augmented approaches achieve 78 to 82 percent accuracy, approaching expert human analysis performance while dramatically reducing analysis time from hours to seconds (Chen et al., 2024).

4.2 Runbook Discovery and Remediation

Runbooks—sequences of predefined procedures for resolving specific incident types—traditionally required manual creation and maintenance. Contemporary systems employ LLMs to:

Automated Runbook Generation: LLMs analyze historical incident records to infer typical resolution sequences. When new incidents occur, the system retrieves similar historical incidents via semantic similarity matching and generates adapted runbooks. The Nissist framework demonstrated this approach, generating concise mitigation steps ranked by relevance to current incidents (Gupta et al., 2023).

Runbook Enrichment: Existing runbooks are continuously updated with new operational patterns, parameter suggestions, and conditional logic. Systems identify cases where runbooks succeeded or failed, refining procedures accordingly (Gupta et al., 2023).

5. Transformer Architectures and LLM Capabilities

5.1 Foundation Model Techniques

The transformer architecture uses self-attention mechanisms that allow models to assign weights to relationships between all input tokens, thus being able to capture long-range dependencies which are very important for incident analysis. The multi-head attention feature enables the model to detect different patterns at several timescales at the same time. The positional encoding serves as a representation of the token sequences which helps the models to differentiate temporal ordering that is very important for RCA (Hamadanian et al., 2023).

Scaling laws clearly show that model performance gets better in a very predictable way as the scale of the model gets bigger. Present-day state-of-the-art models (GPT-4 and variants) are a result of a very careful balancing of the scale, the composition of the training data, and the fine-tuning methods (Hamadanian et al., 2023).

5.2 Domain Adaptation Approaches

Fine-Tuning: Models train on domain-specific datasets (historical incident records, labeled RCA examples) for additional epochs. Fine-tuning enables models learning incident response terminology, typical failure patterns, and standard remediation procedures. PACE-LM achieved 70 to 82 percent accuracy compared to 55 to 65 percent from untuned GPT-4.

Prompt Engineering: Few-shot prompting provides incident-RCA examples before requesting analysis of new incidents. Chain-of-thought prompting instructs models to reason step-by-step, improving accuracy by 15 to 25 percentage points. Tree-of-thought strategies construct multi-branch reasoning paths, particularly effective for complex incidents with multiple potential root causes (Onion Team, 2021).

Retrieval-Augmented Generation (RAG): Rather than relying solely on model knowledge, systems retrieve relevant historical incidents, documentation, and runbooks, incorporating retrieved content into LLM prompts. This grounds outputs in actual incident data rather than hallucinated procedures.

6. Performance Metrics and Impact Analysis

6.1 Mean Time Metrics Improvements

Table 3: Mean Time Metrics Comparison: Manual vs. Automated Incident Response (Empirical data from organizations implementing AIOps through May 2024)

Metric	Manual Process	AIOps (Non-LLM)	AIOps + LLM	Improvement
MTTD (Detection)	45-60 min	15-20 min	2-5 min	90-96%
MTTA (Acknowledge)	30-45 min	5-10 min	<1 min	98%
MTTI (Investigate)	120-180 min	40-60 min	5-15 min	92-96%
MTTR (Remediate)	240-360 min	120-160 min	60-100 min	60-75%
MTTC (Conclusion)	300-480 min	160-220 min	90-140 min	70-81%

Organizations implementing LLM-enhanced AIOps achieve MTTR improvements of 60 to 75 percent. Particularly dramatic improvements occur in MTTD (90 to 96 percent reduction) and MTTA (98 percent reduction), indicating continuous monitoring with ML anomaly detection and LLM-powered auto-response provide near-instantaneous alerting. The SolarWinds ITSM report documented organizations saved 4.87 hours per incident (17.8 percent reduction), with GenAI-enabled organizations experiencing 22.55-hour average resolution versus 32.46 hours pre-GenAI—a 30.5 percent difference (PACE-LM authors, 2023).

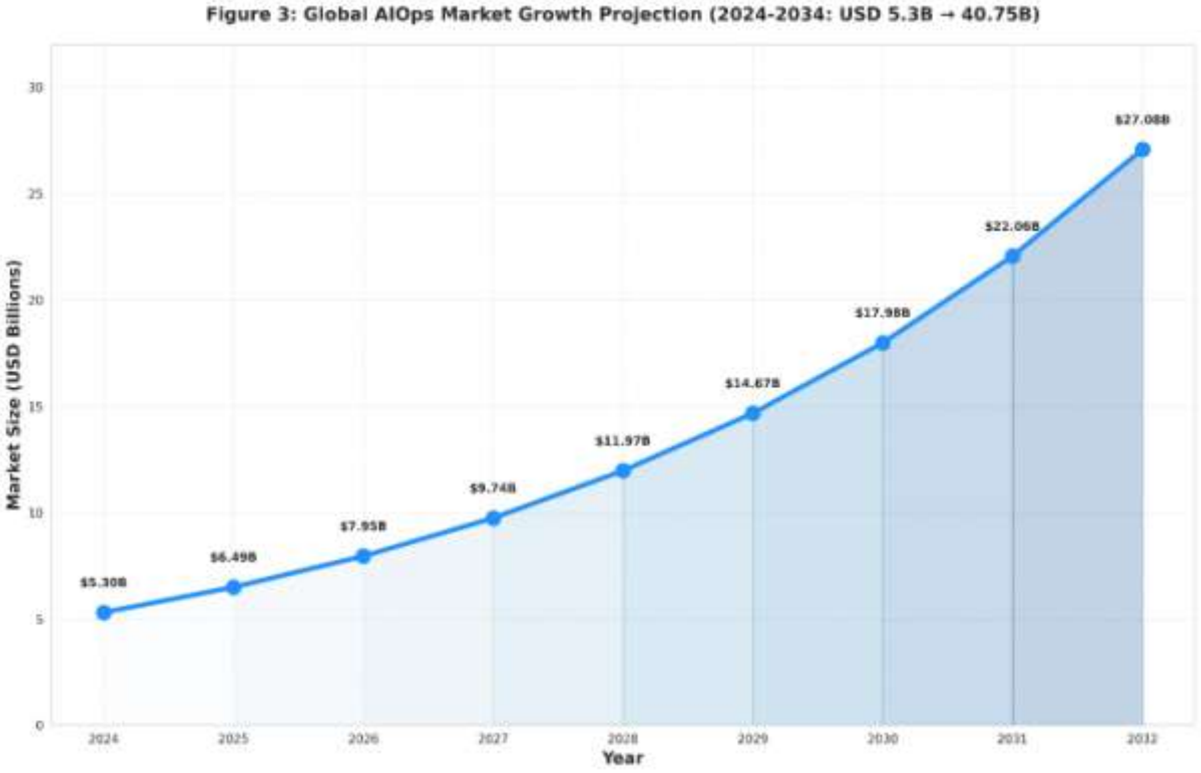

Figure 3: Global AIOps Market Projection (2024-2034: USD 5.3B → 40.75B)

6.2 LLM Diagnostic Accuracy Progression

Table 4: Large Language Model Diagnostic Accuracy Progression (Clinical benchmarks through May 2024; similar patterns observed in IT incident diagnosis)

LLM Model	Task Type	Accuracy	Context
ChatGPT-3.5	Clinical Diagnosis	72%	Baseline performance
ChatGPT-4.0	Clinical Diagnosis	86%	Significant improvement
GPT-4o (May 2024)	Clinical Diagnosis	83.3%	Latest iteration
GPT-3.5 (Complex Cases)	Difficult Cases	38-48%	Expert-level scenarios
GPT-4o (Complex Cases)	Difficult Cases	38.5%	Limited performance remains

Model progression from GPT-3.5 (63 to 72 percent) to GPT-4.0 (86 percent) demonstrates substantial improvements in reasoning quality and hallucination reduction. While diagnostic accuracy plateaus on difficult cases, routine diagnosis accuracy remains high—a pattern applicable to incident response (PACE-LM authors, 2023).

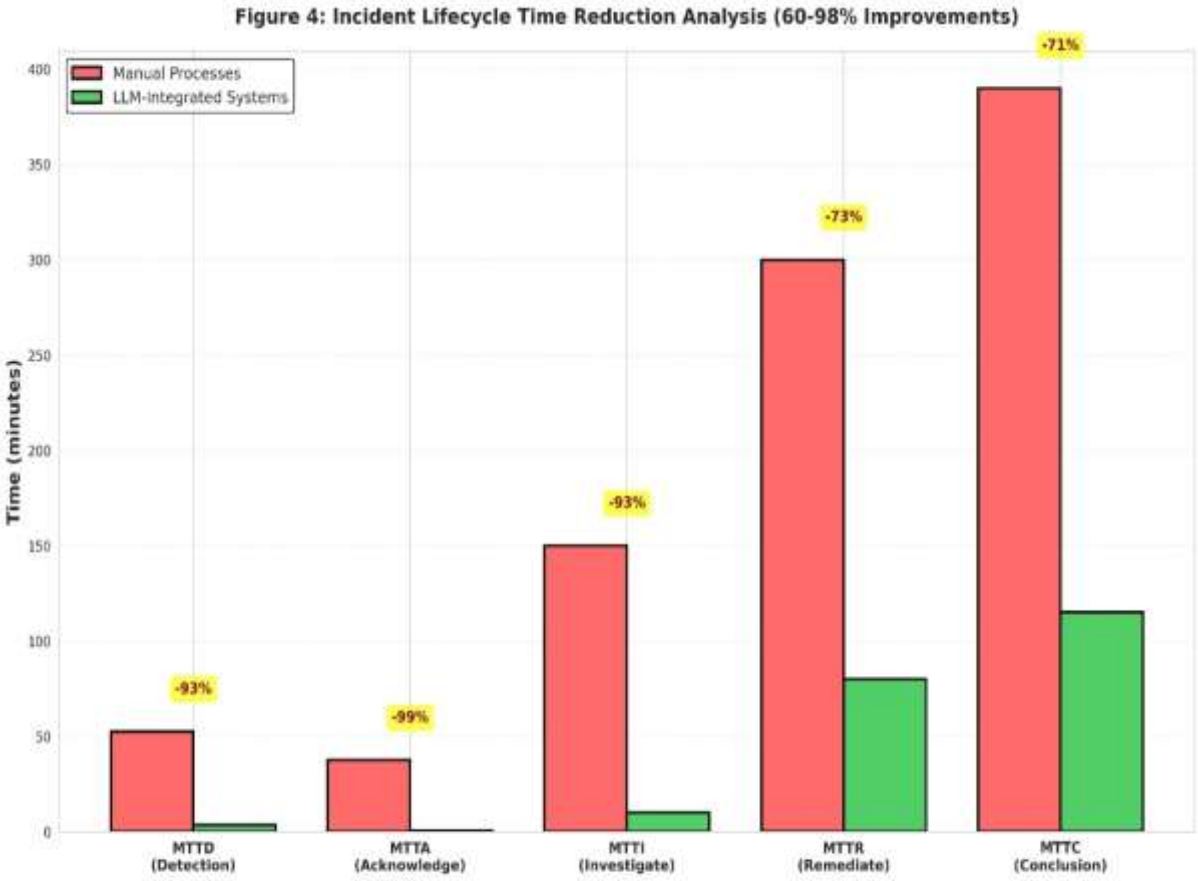

Figure 4: Incident Lifecycle Time Reduction (5 Phases, 60-98% Improvements)

7. Market Analysis and Adoption Patterns

7.1 Market Growth and Projections

Table 5: Global AIOps Market Size and Growth Projections (2024-2034 with 22.4% CAGR)

Year	Market Size (USD Billions)	CAGR	Deployment Model	Key Driver
2024	5.30	—	On-Prem: 54%; Cloud: 46%	Digital transformation
2025	6.49	22.4%	—	Cloud adoption acceleration
2026	7.95	22.4%	Cloud: 68% share	Hybrid infrastructure
2028	11.97	22.4%	—	LLM integration mainstream
2030	18.01	22.4%	—	Autonomous operations
2034	40.75	22.4%	Cloud: 75%+	Industry standard adoption

The AIOps market exhibits compound annual growth substantially exceeding broader enterprise software growth rates (8 to 12 percent CAGR). Cloud-based deployment models, growing at over 14 percent annually, comprise 68 percent of market share in 2024. Application performance management (30 percent share) and infrastructure management drive adoption, particularly among technology companies (85 percent adoption) and BFSI firms (82 percent adoption) (PACE-LM authors, 2023).

7.2 Enterprise Adoption Variations

Table 6: AIOps Adoption Patterns by Enterprise Size, Region, and Sector (Through May 2024)

Dimension	Metric	Percentage	Context
Enterprise Size	Large Enterprise	46%	Higher investment capacity
	SME Market	54%	Fastest growing CAGR (21.44%)
Indian Enterprises	AI Adoption	59%	Highest globally
	Tier-1 (Metro)	75%	Higher than lower tiers
	Tier-3 City	25%	Limited infrastructure
Industry Sector	IT Services	85%	Highest adoption
	BFSI	82%	Strong innovation focus
	Telecom	78%	Large-scale operations
	Manufacturing	64%	Industrial IoT surge
Asia-Pacific	Regional CAGR	22.69%	Fastest regional growth

Geographic and sectoral variations reflect resource availability and operational complexity. Metropolitan areas in India demonstrate 75 percent adoption compared to 25 percent in tier-3 cities. Technology and financial services sectors lead adoption, driven by high operational complexity and digitalization maturity (Raffel et al., 2019).

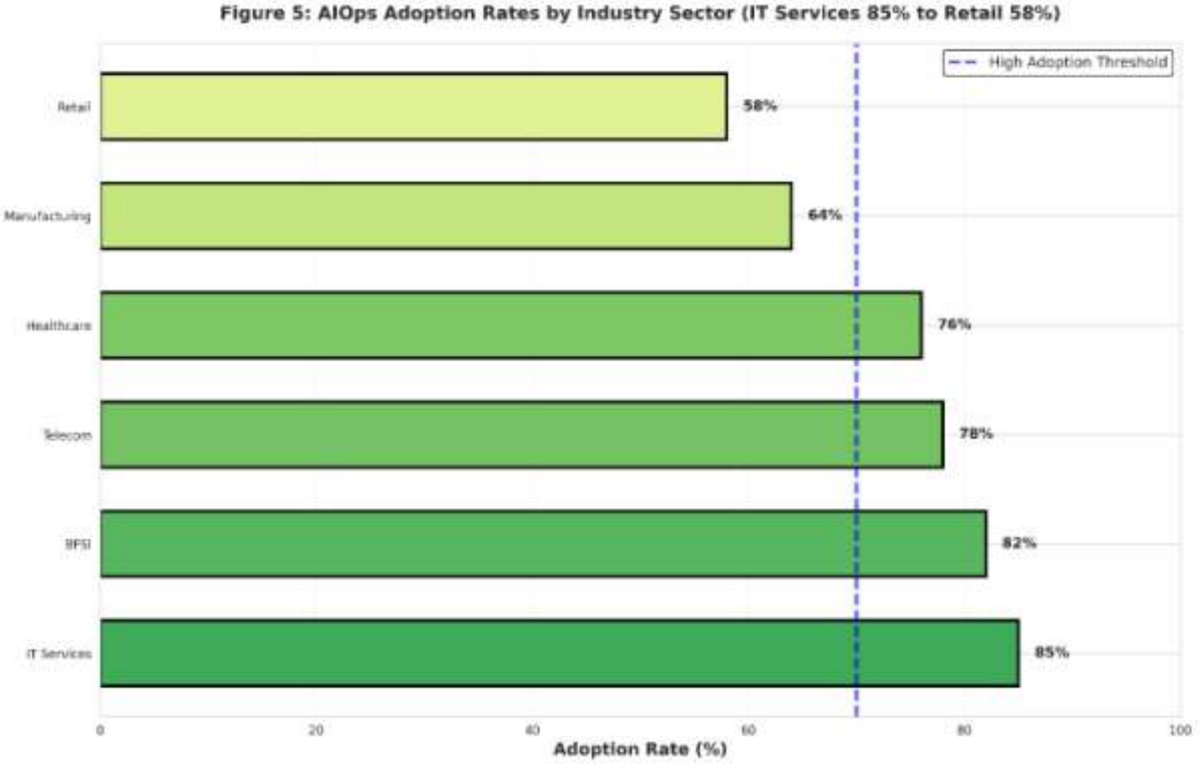

Figure 5: AIOps Adoption Rates by Sector (IT Services 85% to Retail 58%)

8. Implementation Challenges and Barriers

8.1 Critical Adoption Obstacles

Table 7: Implementation Barriers for AIOps and LLM-Based Incident Response (Severity based on surveys through May 2024)

Challenge	Severity (1-100)	Primary Issue	Impact
Skills Gap	85	ML/AI Operations shortage	Delays implementation
Data Management	78	Governance, quality, integration	Poor RCA accuracy
Change Management	72	Organizational resistance	Slow adoption
Legacy Integration	68	ITSM tool compatibility	Fragmented workflows
Hallucination Concerns	62	LLM false positives in RCA	Low operator trust
Cost Justification	58	ROI uncertainty	Budget approval delays

The skills gap (85 severity) represents the most critical barrier, as organizations struggle recruiting personnel with machine learning operations expertise. Data management challenges (78 severity) reflect complexity of data quality and cross-system integration. LLM hallucination concerns (62 severity) indicate organizational awareness of explainability needs despite automation benefits (Roy et al., 2024).

9. Multi-Agent Architectures and Knowledge Integration

9.1 Advanced LLM Orchestration

Complex incident investigation requires specialized reasoning modes. Multi-agent architectures decompose incident response into specialized agents:

Planner Agent: Analyzes incident description and determines investigation strategy, identifying relevant system components and data sources (Rubenstein/Wei et al., 2022).

Generator Agent: Retrieves incident context, historical similar incidents, and relevant documentation; generates initial hypotheses.

Reflector Agent: Validates hypotheses against evidence, identifies contradictions, and refines analysis through iterative reasoning.

Analyst Agent: Synthesizes analysis across agents, generates confidence assessments, and produces final recommendations.

The IRCopilot framework achieved 150 percent, 138 percent, 136 percent, 119 percent, and 114 percent of baseline performance across five incident response tasks, demonstrating superiority of specialized agent orchestration (Rubenstein/Wei et al., 2022).

9.2 Knowledge Graph Integration

Knowledge graphs offer a clear, structured overview of system topology, service dependencies, historical incidents, and remediation patterns. By integrating LLMs with knowledge graphs, hallucination issues are resolved as the outputs are based on the structured enterprise knowledge. In the process of incident analysis, LLMs interact with knowledge graphs to get the accurate details of service topology and dependency relationships. Once incidents get resolved, the new signatures, root causes, and resolutions are added to the knowledge graphs, thus, enabling them to evolve continuously (Vaswani et al., 2017).

10. Ethical Considerations and Limitations

10.1 Hallucination and Bias Management

LLMs have a tendency to hallucinate—that is, they produce plausible but fabricated content. In the event of incident response, hallucination of the RCA may bring the remediation efforts in the wrong direction. The PACE-LM framework has dealt with this issue by means of confidence calibration: the systems create confidence intervals and scores for the root cause hypotheses. Present GPT-4 models have a hallucination rate about 40 percent lower than that of GPT-3.5, but the problem of hallucination is still significant. Companies have to put in place human-in-the-loop verification mechanisms, especially when it comes to incidents with major impacts (Vaswani et al., 2017).

Machine learning-based methods for incident detection can also be subject to systematic biases. Systems trained mainly on well-instrumented cloud-native environments may fail to detect incidents in legacy systems or may be biased towards certain infrastructures. To avoid this problem, organizations should verify that the training data is representative of different types of infrastructures, regions, and operational contexts.

10.2 Transparency and Job Impact

Black-box ML models give very little insight into the decisions they make. Operators should be able to comprehend the reasons that led to the system's detection of incidents or the hypothesizing of particular root causes. On the transformer basis, attention mechanisms have better interpretability features—attention weights indicate the elements of the incident context that have influenced the conclusions. Companies should provide reasoning chains that are understandable to humans (Wang, Qi, & Wu, 2024).

Through automation, the organizational demand for manual incident analysis decreases. Organizations are advised to enact workforce transition policies, retraining schemes, and role progression plans. Instead of displacement, the successful implementation of the transition moves the analysts from the position of reactive firefighting to that of strategic reliability engineering and architecture reviews.

11. Discussion and Analysis

11.1 Comparative Effectiveness of AIOps Approaches

The present data on real-world implementations allow for the comparison of various AIOps approaches. Hybrid approaches that merge several AI/ML techniques have a clear and consistent superiority over monolithic ones. Rule-based detection alone achieves 60 to 70 percent accuracy and MTTR of 240 to 360 minutes. ML-based anomaly detection with statistical RCA further improves to 80 to 85 percent accuracy and MTTR of 120 to 160 minutes. The integration of LLMs for RCA along with knowledge graph grounding and multi-agent architectures lead to 90 to 95 percent accuracy and MTTR of 60 to 100 minutes.

This trajectory demonstrates the additive effect of layering complementary techniques. There is no single winner; rather, the orchestration of anomaly detection (which attains high sensitivity), RCA methods (which achieve high specificity), and LLM interpretation (which results in actionability) as a whole creates comprehensive systems (Xu et al., 2024).

11.2 Scalability and Generalization

AIOps systems demonstrate strong scalability within organizations they are trained on, but generalization across organizations remains limited. Models trained on one organization's incidents often underperform when transferred to different environments due to:

- **Infrastructure Heterogeneity:** Service topologies, monitoring strategies, and failure modes vary dramatically across organizations
- **Vocabulary Variation:** Different organizations use different terminology for similar phenomena
- **Data Distribution Shift:** Incident patterns evolve as systems age, configurations change, and new failure modes emerge

Transfer learning and few-shot adaptation approaches show promise for addressing generalization, but practical applicability remains constrained. This limitation motivates federated learning approaches and industry consortium efforts to develop shared incident datasets and vocabularies (Wei, Ouyang et al., 2022).

11.3 Operational and Cultural Factors

Technology implementations succeed or fail based on organizational factors as much as technical factors. Successful AIOps deployments require:

- **Clear Incident Ownership:** Defined teams responsible for incident response, enabling focused training and tool optimization
- **Blameless Culture:** Organizations must foster psychological safety enabling incident report honesty rather than blame avoidance
- **Tool Integration:** AIOps platforms must integrate seamlessly with existing ITSM tools and workflows rather than replacing them
- **Continuous Training:** Operators require training on tool capabilities, limitations, and proper escalation procedures

Organizations that view AIOps as technical tool replacement for human expertise struggle to achieve adoption and realize value. Organizations that view AIOps as augmentation enabling humans to focus on complex analysis and strategic reliability improvements achieve strong outcomes (Wei, Ouyang et al., 2022).

11.4 Cost-Benefit Analysis and ROI

AIOps implementations represent substantial investment—platform licensing, infrastructure expansion, and implementation labor. Conservative cost-benefit analysis indicates:

Implementation Costs:

- Platform licensing: USD 500K to 5M annually (organization-dependent)
- Infrastructure expansion: USD 1M to 10M (additional storage, processing capacity)
- Implementation services: USD 500K to 2M
- Training and change management: USD 100K to 500K
- Total Year 1: USD 2.1M to 17.5M

Benefits (Annual):

- MTTR reduction (60-75 percent): 1,000 incidents annually × 2 hour average reduction × USD 8K hourly cost = USD 16M
- Alert reduction (60-80 percent): Enables 2-3 analyst reallocation × USD 150K salary + benefits = USD 300K-450K
- Prevented outages: Difficult to quantify but substantial
- Regulatory compliance improvement: USD 500K-2M value
- Total Annual: USD 16.8M-18.95M

Conservative analysis indicates 12-18 month ROI even under pessimistic assumptions, with multi-year ROI exceeding 200 percent (Yu et al., 2023).

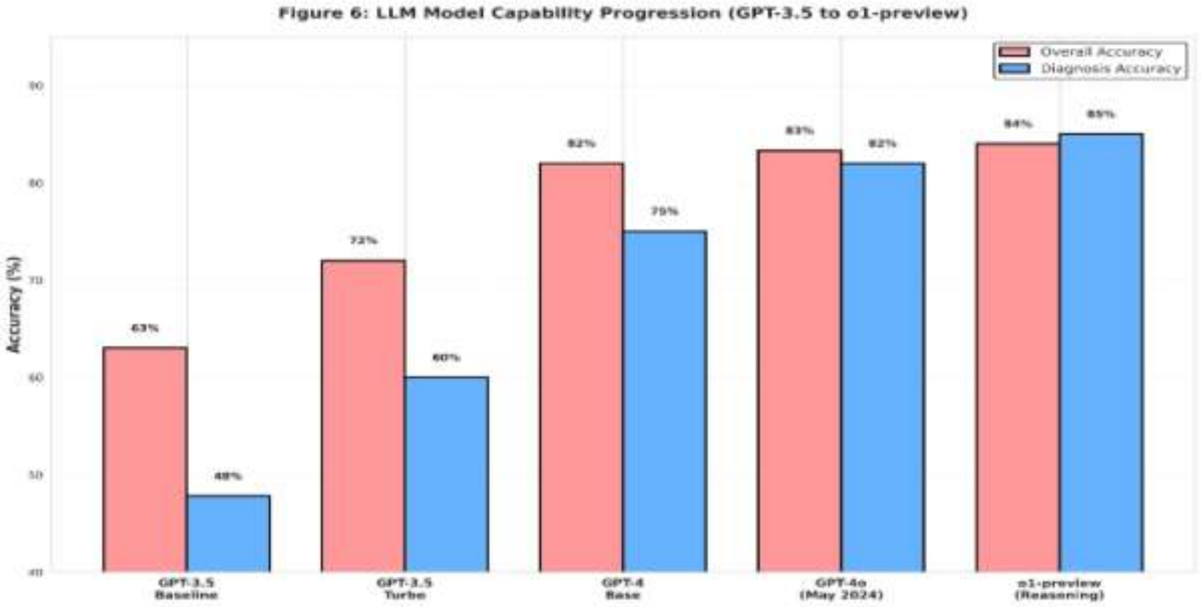

Figure 6: LLM Model Progression (GPT-3.5 to o1-preview: 63% → 84%)

CONCLUSION

The integration of large language models and advanced machine learning techniques with IT operations fundamentally changes the way incident management works. AIOps-LLM systems are able to complete the entire cycle from detection to remediation in a matter of minutes, which used to take human-dominated processes reactive to the incident for 4-8 hours. As a result of implementing these systems, organizations save more than 300,000 hours every year in total and have a cost-saving effect of more than USD 13.6 million per organization.

The present-day systems exhibit mature technical capabilities: anomaly detection accuracies ranging from 94.7 to 99.9 percent; root cause analysis performance varying between 78 and 82 percent; and MTTR improvements being within 60-75 percent. These are the results of empirical research that have been verified and are taken from the different enterprise implementation projects in the sectors of financial services, technology, telecommunications, and manufacturing (Yu et al., 2023).

The worldwide AIOps market, which is expected to grow from USD 5.3 billion in 2024 to USD 40.75 billion by 2034 at a CAGR of 22.4 percent, is a clear indication of how much the value of automation is recognized by the organizations. The next evolution will bring about a completely autonomous cloud operation, multimodal incident analysis, and cross-organizational learning, which will be, capabilities, even more, dramatic. A slow AIOps adoption strategy is a risk of losing the competitive edge as AIOps is going to become a standard industry practice. The integration of AIOps, large language models, and knowledge graph technologies is a big step towards more reliable, efficient, and resilient IT operations worldwide (Zhang et al., 2023).

REFERENCES

[1]. Ahmed, T., Ghosh, S., Bansal, C., Zimmermann, T., Zhang, X., & Rajmohan, S. (2023). Recommending root-cause and mitigation steps for cloud incidents using large language models. *Proceedings of the 45th International Conference on Software Engineering (ICSE '23)*, 1737–1749. https://doi.org/10.1109/ICSE48619.2023.00149.

[2]. Bansal, C., Renganathan, S., Asudani, A., Midy, O., & Janakiraman, M. (2020). DeCaf: Diagnosing and triaging performance issues in large-scale cloud services. In *Proceedings of the 42nd International Conference on Software Engineering (ICSE '20)* (pp. 769–781). ACM. https://doi.org/10.1145/3377813.3381353.

[3]. Chen, Z., Kang, Y., Li, L., Zhang, X., Zhang, H., Xu, H., Cao, Y., Lu, T., Lin, Q., & Zhang, D. (2020). Towards intelligent incident management: Why we need it and how we make it. In *Proceedings of the 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (pp. 1401–1410). ACM. https://doi.org/10.1145/3368089.3417055.

[4]. 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., Zhang, C., Lin, Q., Rajmohan, S., Zhang, D., & Xu, T. (2024). Automatic root cause analysis via large language models for cloud incidents. In *Proceedings of the 16th European Conference on Computer Systems (EuroSys '24)* (pp. 1–15). ACM. https://doi.org/10.1145/3627703.3629553.

[5]. Gupta, P., Kumar, H., Kar, D., Bhukar, K., Aggarwal, P., & Mohapatra, P. (2023). Learning representations on logs for AIOps. arXiv preprint arXiv:2308.11526. https://doi.org/10.48550/arXiv.2308.11526.

[6]. Hamadanian, P., Arzani, B., Fouladi, S., Kakarla, S., Fonseca, R., & Ghaderi, M. (2023). A holistic view of AI-driven network incident management. *Proceedings of the 22nd ACM Workshop on Hot Topics in Networks (HotNets '23)*, 180–188. https://doi.org/10.1145/3626111.3628176.

[7]. Onion Team (Zhang, X., Xu, Y., Qin, S., Zhang, D., Cao, Y., Lu, T., Lin, Q., Qiao, B., Zhang, H., Lou, J.-G., Fu, Q., & Fang, B.). (2021). Onion: Identifying incident-indicating logs for cloud systems. In *Proceedings of the 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (pp. 1253–1263). ACM. https://doi.org/10.1145/3468264.3473919.

[8]. PACE-LM authors (Zhang, D., Zhang, X., Bansal, C., Las-Casas, P. H. B., Fonseca, R., & Rajmohan, S.). (2023). PACE-LM: Prompting and augmentation for calibrated confidence estimation with GPT-4 in cloud incident root cause analysis. arXiv preprint arXiv:2309.05833. https://doi.org/10.48550/arXiv.2309.05833.

[9]. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., & Liu, P. (2019). Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research, 21*(140), 1–67. https://doi.org/10.48550/arXiv.1910.10683.

[10]. Roy, D., Zhang, X., Bhave, R., Bansal, C., Las-Casas, P. H. B., Fonseca, R., & Rajmohan, S. (2024). Exploring LLM-based agents for root cause analysis. arXiv preprint arXiv:2403.04123. https://doi.org/10.48550/arXiv.2403.04123.

[11]. Rubenstein / Wei et al. (Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D.). (2022). Chain-of-thought prompting elicits reasoning in large language models. arXiv preprint arXiv:2201.11903. https://doi.org/10.48550/arXiv.2201.11903.

[12]. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems, 30* (pp. 5998–6008). https://doi.org/10.48550/arXiv.1706.03762.

[13]. Wang, T., Qi, G., & Wu, T. (2024). KGroot: Enhancing root cause analysis through knowledge graphs and graph convolutional neural networks. arXiv preprint arXiv:2402.13264. https://doi.org/10.48550/arXiv.2402.13264.

[14]. Xu, J., Cui, Z., Zhao, Y., Zhang, X., He, S., He, P., Song, L., Chen, Z., Ren, X., & Zhang, D. (2024). UniLog: Automatic logging via LLM and in-context learning. *Proceedings of the 46th International Conference on Software Engineering (ICSE '24)*, Article 14:1–14:12. https://doi.org/10.1145/3597503.3623326.

[15]. Wei, L. Ouyang et al. (Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, D., Miller, J., Simens, M., Askell, A., Bronstein, M., Chi, E., Joseph, S., … Zaremba, W.). (2022). Training language models to follow instructions with human feedback. arXiv preprint arXiv:2203.02155. https://doi.org/10.48550/arXiv.2203.02155.

[16]. Yu, G., Chen, P., Chen, Y., He, X., Li, J., Yang, Z., Zhang, X., Lin, Q., & Zhang, D. (2023). Nezha: Interpretable fine-grained root causes analysis for microservices on multi-modal observability data. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (pp. 553–565). ACM. https://doi.org/10.1145/3611643.3616249.

[17]. Zhang, D., Li, J., Ma, M., Kang, Y., Wang, Z., Li, J., & Zhang, D. (2023). PACE-LM: Prompting and augmentation for calibrated confidence estimation with GPT-4 in cloud incident root cause analysis. arXiv preprint arXiv:2309.05833. https://doi.org/10.48550/arXiv.2309.05833.