IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

ISO 3297:2007 Certified ⁜ Impact Factor 7.12 ⁜ Vol. 9, Issue 11, November 2022

DOI: 10.17148/IARJSET.2022.91122

ENABLING INTELLIGENT INFRASTRUCTURE:AI-DRIVEN AUTOMATION AND RESILIENCE IN CLOUD-NATIVE SYSTEMS

Prudhvi Naayini[1], Srikanth Kamatala[2]

Independent Researcher, Dallas, TX, USA[1]

Independent Researcher Dallas, TX, USA[2]

Abstract: Cloud native infrastructure management is being transformed by Artificial Intelligence (AI) and Machine Learning (ML) techniques, often referred to as AIOps, which automate complex operations and enhance system resilience. AIOps capabilities encompass predictive maintenance, forecasting and preventing failures before they impact services, intelligent observability through the analysis of logs, metrics, and traces, and autonomous fault remediation that enables self-healing systems. These approaches are particularly valuable in Kubernetes based architectures, where dynamic microservices environments generate massive volumes of telemetry data that AI can analyze to proactively detect anomalies and performance issues.

Major cloud platforms have integrated AI driven automation into their operations toolchains. For instance, AWS DevOps Guru employs ML models to identify operational anomalies and recommend remediation actions, while Azure Monitor and Google Cloud Operations embed machine learning for intelligent alerting, performance tuning, and capacity forecasting. Open source and hybrid tools further enrich this ecosystem. KubeFlow supports ML workflows on Kubernetes, and observability frameworks like Prometheus and Elastic APM collect telemetry data that feeds into AI driven analytics and automated responses.

This article highlights how AI driven automation and AIOps practices are enhancing infrastructure reliability and efficiency, while also addressing persistent challenges. These include model drift where model accuracy degrades as systems evolve, poor data quality that undermines analytical insights, and a lack of explainability in AI decisions which complicates trust and broader adoption of AIOps solutions.

Keywords: AI-Driven Infrastructure, AIOps, Cloud-Native Systems, Predictive Maintenance

I. INTRODUCTION

Cloud-native architectures, consisting of containerized applications, microservices, and orchestration platforms such as Kubernetes, have redefined scalability and agility in modern digital ecosystems [1, 2]. However, these advancements have also introduced operational complexity that challenges the capabilities of traditional, rule-based infrastructure management approaches. The distributed and dynamic nature of such environments demands intelligent automation and adaptive system behavior.

Artificial Intelligence for IT Operations (AIOps) has emerged as a promising solution to address these complexities. By leveraging AI and Machine Learning (ML), AIOps enables predictive analytics, automated diagnostics, and autonomous remediation [3]. These technologies support a paradigm shift from reactive troubleshooting to proactive, data-driven infrastructure management [4].

Kubernetes environments, in particular, benefit from AIOps techniques. Due to the ephemeral and distributed nature of microservices, such platforms generate extensive telemetry data in the form of logs, metrics, and traces. AI models can process this data to identify anomalies, optimize autoscaling, and support intelligent workload placement [5]. Industry-leading cloud platforms, including AWS DevOps Guru [6], Azure Monitor, and Google Cloud Operations Suite [7], have integrated AIOps capabilities for anomaly detection, performance optimization, and capacity forecasting.

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

ISO 3297:2007 Certified ⁜ Impact Factor 7.12 ⁜ Vol. 9, Issue 11, November 2022

DOI: 10.17148/IARJSET.2022.91122

Additionally, open-source tools such as Kube Flow, Prometheus, and Elastic APM support the development of AI-powered observability pipelines, further promoting the integration of AIOps into hybrid and multi-cloud deployments [8]. Despite these advancements, key challenges persist. Issues such as model drift [5], inconsistent data quality [3], and limited explainability in AI decisions continue to impede widespread adoption. Retrofitting AIOps into legacy systems and ensuring transparency and accountability in automated operations remain critical concerns.

This paper explores the architecture, design principles, and real-world implementations of AI-driven infrastructure automation. It highlights the transformative potential of AIOps in enhancing infrastructure resilience while critically evaluating the limitations and considerations that influence adoption in production environments.

II. BACKGROUND AND RELATED WORK

The convergence of Artificial Intelligence (AI) and infrastructure management has emerged as a critical response to the operational complexity of modern digital systems. As infrastructures transitioned from monolithic to distributed and cloud-native architectures, traditional rule-based monitoring and manual troubleshooting proved inadequate [1]. The need for scalable, adaptive, and intelligent operational mechanisms gave rise to AIOps Artificial Intelligence for IT Operations.

AIOps applies machine learning, statistical modeling, and pattern recognition to ingest and analyze vast volumes of IT telemetry data such as logs, metrics, traces, and events [4]. These systems automate anomaly detection, incident correlation, root cause identification, and even real-time remediation [3]. In microservices-based environments, where service interdependencies are highly dynamic, AIOps significantly reduces the mean time to detect (MTTD) and mean time to resolution (MTTR).

Several studies have demonstrated the feasibility and value of AI-enhanced infrastructure. Katsikas et al. [8] proposed a framework for self-healing distributed systems using anomaly detection and AI-based policy engines. Similarly, Tuli et al. [5] illustrated how deep learning models deployed at the edge can support predictive maintenance and fault prevention, offering a blueprint for AI-driven resilience in hybrid architectures.

The widespread adoption of Kubernetes has accelerated the integration of AIOps into production systems. Native observability tools such as Prometheus and Elastic APM enable high-resolution telemetry capture, while platforms like KubeFlow orchestrate end-to-end machine learning workflows within Kubernetes environments [2]. These tools facilitate the implementation of intelligent automation pipelines capable of learning and adapting continuously.

Leading cloud service providers have recognized this paradigm shift. AWS DevOps Guru uses ML to detect anomalies and recommend remediation steps [6], while Google Cloud Operations Suite and Azure Monitor embed predictive analytics to enhance performance and availability [7]. These platforms represent the commercialization of AIOps concepts initially developed in academic and open-source contexts.

Despite the promise of AIOps, challenges remain. Model drift can degrade prediction accuracy over time, while low-quality data can lead to misleading inferences [5]. Furthermore, the lack of transparency in AI decision-making raises concerns about trust, compliance, and accountability. Ongoing research is exploring hybrid approaches that combine traditional rule-based systems with adaptive AI to strike a balance between automation and human oversight.

III. AI-DRIVEN AUTOMATION LAYERS

AI-driven infrastructure management operates through a multi-layered pipeline that systematically processes telemetry data to deliver intelligent, automated operational capabilities. These layers—ranging from data collection to autonomous remediation—work together to enable proactive and resilient system behavior in cloud-native environments.

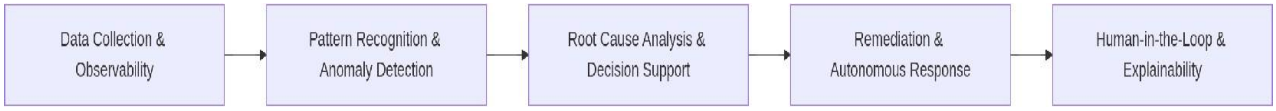

Figure 1: AI-Driven Automation Layers in Cloud-Native Systems

1.1 Data Collection and Observability

The foundation of AIOps lies in comprehensive observability. Cloud-native systems emit vast volumes of telemetry data, including logs, metrics, distributed traces, and event streams. Tools like Prometheus and Elastic APM enable the collection of fine-grained monitoring data in Kubernetes environments [2].

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

ISO 3297:2007 Certified ⁜ Impact Factor 7.12 ⁜ Vol. 9, Issue 11, November 2022

DOI: 10.17148/IARJSET.2022.91122

This telemetry is typically stored in time-series databases or log aggregation systems, providing the raw input required for downstream machine learning analytics.

1.2 Pattern Recognition and Anomaly Detection

Once collected, telemetry data is analyzed using machine learning models to detect deviations from normal behavior. These models based on supervised, unsupervised, or semi-supervised learning can identify subtle performance anomalies and anticipate system failures [3]. For example, AWS DevOps Guru applies ML algorithms to identify operational anomalies, correlate system events, and generate remediation insights [6]. Such models are especially useful in dynamic microservices environments, where traditional threshold-based monitoring is insufficient.

1.3 Root Cause Analysis and Decision Support

AIOps platforms extend beyond anomaly detection to support decision-making and root cause analysis. By correlating disparate events and system behaviors, these platforms identify the origin of performance degradations and recommend targeted interventions. Techniques such as clustering, dependency graph modeling, and causal inference are employed to connect symptoms to systemic faults [8]. This is essential in distributed architectures, where failures often propagate through multiple services.

1.4 Remediation and Autonomous Response

Insights generated from AI models can be used to trigger remediation actions, ranging from alerting operators to fully automated responses. Common actions include restarting failed pods, reallocating workloads, scaling resources, or triggering CI/CD pipelines. Kubernetes provides native support for such orchestration, enabling seamless integration between AIOps insights and automated infrastructure response [5]. In production environments, policy-driven controls are often applied to limit autonomous actions to safe, approved boundaries.

1.5 Human-in-the-Loop and Explainability

While automation is a central goal of AIOps, human oversight remains crucial, particularly in mission-critical systems. Explainable AI (XAI) techniques provide interpretability into the decision-making process, helping engineers understand model outputs and maintain trust [3]. Human-in-the-loop design ensures that operators can validate AI-driven actions and override or adjust behavior based on domain knowledge and context.

IV. PROPOSED ARCHITECTURE

To operationalize AI-driven automation in cloud-native environments, we propose a modular architecture composed of five tightly integrated layers. These layers correspond to the automation pipeline previously discussed and are designed to support observability, intelligence, remediation, and governance across distributed systems. The architecture leverages Kubernetes as the orchestration backbone and integrates open-source and cloud-native tools for scalability and resilience.

1.6 Layer 1: Telemetry and Observability

This foundational layer is responsible for collecting real-time telemetry data from infrastructure components, containers, services, and applications. It integrates tools such as Prometheus for metrics collection, Elastic APM for tracing, and Fluentd or Logstash for log aggregation [2]. Data is normalized and sent to time-series databases or analytics pipelines for further processing.

1.7 Layer 2: Data Processing and Feature Extraction

Once ingested, raw telemetry data is transformed and enriched. Feature extraction methods include statistical summarization, temporal windowing, and embedding generation. This layer can be implemented using stream processing frameworks (e.g., Apache Flink) or Kubernetes-native tools such as KNative or Kafka on Kubernetes [1]. Extracted features serve as input for anomaly detection models and system behavior profiling.

1.8 Layer 3: AI/ML Inference Engine

The core intelligence layer hosts pre-trained and online-learning models for anomaly detection, classification, and root cause prediction. These models can be served using KubeFlow pipelines, TensorFlow Serving, or ONNX runtimes deployed as Kubernetes services [5]. The inference engine communicates bidirectionally with observability systems and feeds into alerting and decision-making components.

1.9 Layer 4: Automation and Orchestration Engine

This layer connects insights from the inference engine to orchestrated actions. It leverages Kubernetes-native automation

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

ISO 3297:2007 Certified ⁜ Impact Factor 7.12 ⁜ Vol. 9, Issue 11, November 2022

DOI: 10.17148/IARJSET.2022.91122

mechanisms such as Horizontal Pod Autoscalers (HPA), Kubernetes Operators, and custom controllers to take predefined or AI-suggested remediation actions [8]. Integration with CI/CD systems (e.g., ArgoCD, JenkinsX) allows for continuous adaptation based on operational intelligence.

1.10 Layer 5: Policy and Explainability Interface

To ensure human oversight and accountability, this top layer exposes interfaces for policy configuration, explainable AI (XAI), and human-in-the-loop validation. It allows engineers to approve, audit, or override AI-driven decisions and provides transparency into the rationale behind model outputs [3]. This layer also logs AI actions and supports compliance monitoring in regulated environments.

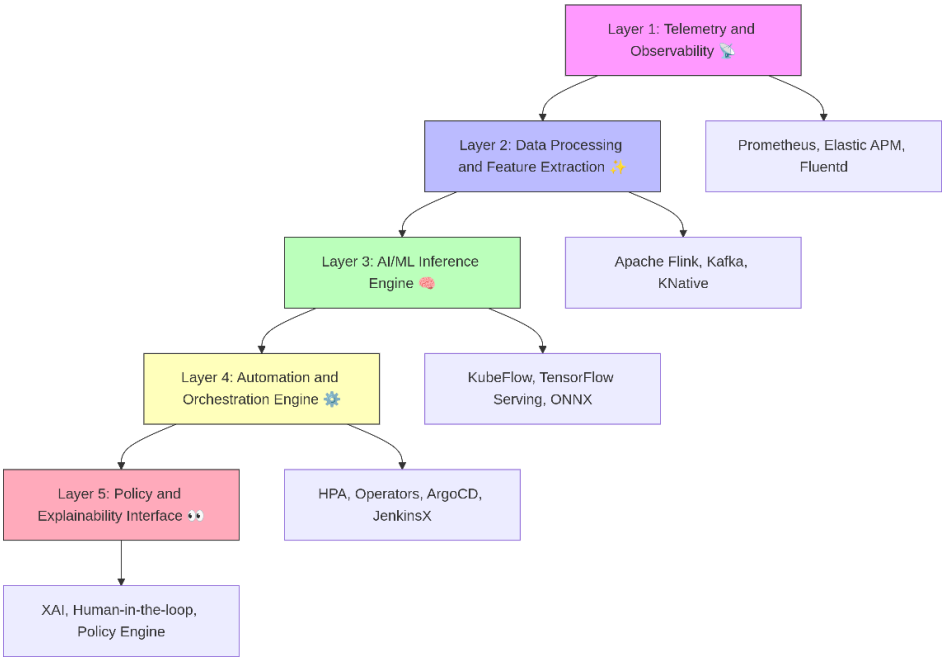

Figure 2: Proposed Five-Layer AI-Driven Infrastructure Architecture

1.11 Deployment Considerations

The architecture is cloud-agnostic and supports hybrid deployments across public cloud, private data centers, and edge clusters. Containerized microservices and infrastructure-as-code (IaC) tools such as Helm and Terraform are used for reproducibility and scalability. Security and governance are enforced using Role-Based Access Control (RBAC), mutual TLS, and secure model delivery pipelines.

V. CASE STUDIES AND REAL-WORLD APPLICATIONS

The practical implementation of AI-driven infrastructure automation has gained significant traction across both commercial cloud platforms and open-source ecosystems. This section highlights key case studies that demonstrate how AIOps architectures are deployed to enhance observability, predictive maintenance, and autonomous remediation.

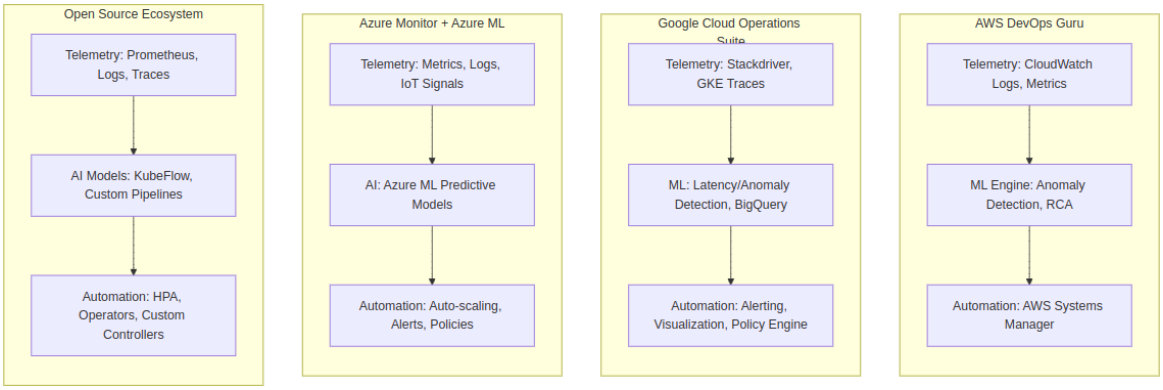

Figure 3: AIOps Pipeline Across Major Platforms and Open-Source Ecosystem

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

ISO 3297:2007 Certified ⁜ Impact Factor 7.12 ⁜ Vol. 9, Issue 11, November 2022

DOI: 10.17148/IARJSET.2022.91122

1.12 AWS DevOps Guru

Amazon Web Services introduced DevOps Guru as a managed AIOps service that applies machine learning models to operational telemetry. The platform ingests logs, metrics, and traces from AWS services such as EC2, RDS, Lambda, and ECS. By continuously analyzing this data, DevOps Guru identifies anomalies, surfaces potential root causes, and recommends corrective actions [6]. It integrates with AWS Systems Manager for automated remediation and provides visual insights through CloudWatch dashboards.

1.13 Google Cloud Operations Suite

Google Cloud's Operations Suite (formerly Stackdriver) leverages AI to enhance observability, alerting, and performance optimization. The suite integrates metrics, logs, and traces from cloud-native services such as GKE (Google Kubernetes Engine) and Cloud Run [7]. Machine learning models detect latency spikes, error anomalies, and resource bottlenecks, enabling proactive incident response. Integration with BigQuery also supports long-term trend analysis and model training on historical data.

1.14 Azure Monitor and Azure Machine Learning

Microsoft Azure integrates AIOps across Azure Monitor and Azure Machine Learning platforms. Azure Monitor collects telemetry across compute, networking, and storage resources, while AI capabilities such as anomaly detection, auto-tuning, and autoscaling are embedded within resource groups [5]. Azure Machine Learning allows for the deployment of predictive maintenance models that act on signals from Azure IoT and log analytics, enabling hybrid infrastructure monitoring.

1.15 Open Source Kubernetes Ecosystem

In open-source environments, Kubernetes clusters augmented with Prometheus, Grafana, and KubeFlow represent a powerful AIOps platform. Prometheus scrapes metrics from pods and nodes, while Grafana dashboards visualize trends and alerts. KubeFlow orchestrates ML pipelines that process observability data for anomaly detection and predictive forecasting [2]. Kubernetes Operators and custom controllers are used to trigger auto-remediation actions, such as restarting pods or scaling deployments.

1.16 Industry Use Cases

Telecom: Edge-based Kubernetes clusters in telecom networks employ AI for proactive node monitoring and traffic rerouting. **Finance:** AIOps supports fraud detection and infrastructure security by identifying unusual traffic and triggering automated isolation protocols. **Manufacturing:** Predictive maintenance models deployed at the edge reduce downtime and optimize production line throughput using streaming telemetry data [5].

These real-world implementations validate the architecture proposed in Section 4, demonstrating that AI-driven automation is not only feasible but increasingly essential for achieving resilient and self-healing infrastructure.

VI. CHALLENGES AND LIMITATIONS

While AI-driven infrastructure automation holds significant promise, its broader adoption faces several technical and operational challenges that must be addressed to ensure robustness, trust, and scalability.

One of the foremost challenges is model drift, where the accuracy and reliability of machine learning models degrade over time as system workloads, configurations, and behavior patterns evolve. In dynamic cloud-native environments, this drift can result in false positives or missed anomalies [5]. To counter this, organizations must adopt continuous retraining strategies and incorporate model monitoring pipelines that can detect and adapt to shifting conditions.

Closely related is the issue of data quality and diversity. AIOps systems depend heavily on structured and timely telemetry data gathered from various sources, including public cloud, edge devices, and legacy infrastructure. Incomplete, inconsistent, or noisy data can lead to flawed insights and unreliable predictions [3]. Enforcing telemetry standards, deploying data validation layers, and improving schema management are crucial to maintaining high-quality inputs.

Another key limitation is the lack of explainability in AI models. Many deep learning or ensemble techniques operate as black boxes, making it difficult for engineers to understand or trust the rationale behind automated decisions. This is especially critical in safety-sensitive or regulated environments where auditability and justification are mandatory. The integration of explainable AI techniques, including interpretable models and visualization tools, is necessary to improve transparency and build operator confidence [4].

Integration with legacy systems also presents substantial hurdles. While modern microservices are designed to emit rich observability signals, many legacy systems lack native support for metrics, logs, or traces in standard formats.

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

ISO 3297:2007 Certified ⁜ Impact Factor 7.12 ⁜ Vol. 9, Issue 11, November 2022

DOI: 10.17148/IARJSET.2022.91122

Bridging these systems requires the development of custom adapters or telemetry gateways, which increases deployment complexity and slows time to value [8].

Furthermore, the cost and resource overhead of running AI workloads is nontrivial. Real-time inference, anomaly detection, and online learning demand compute and storage resources that can strain budgets, especially in edge or hybrid environments [5]. Balancing the trade-off between inference performance and infrastructure efficiency is an ongoing concern.

Finally, while automation is the ultimate goal of AIOps, human oversight remains essential. In mission-critical environments, organizations must retain control over AI-driven actions through approval workflows, policy enforcement, and override mechanisms. Designing effective human-in-the-loop systems ensures accountability, safety, and operational alignment.

These challenges underscore the need for careful architectural design, iterative validation, and a balanced approach that integrates both intelligent automation and human governance.

VII. FUTURE OUTLOOK AND RESEARCH DIRECTIONS

The future of AI-driven operations in cloud-native infrastructure points toward increasingly intelligent, adaptive, and secure systems. As organizations seek to reduce operational complexity while maintaining resilience and compliance, several research directions are expected to shape the next generation of AIOps platforms.

A primary area of innovation is the development of hybrid reasoning systems that integrate rule-based logic with machine learning and neural models. These approaches seek to combine the transparency of symbolic reasoning with the flexibility and generalization capacity of statistical methods [9]. Such architectures offer promise in domains that demand compliance, traceability, and human-in-the-loop validation, enabling more explainable and context-aware automation.

Table 1: Key Challenges in AIOps Adoption and Mitigation Strategies

Challenge	Implications	Mitigation Strategy
Model Drift and Adaptability	Reduced accuracy over time; false positives or missed anomalies	Implement continuous monitoring and retraining pipelines; use drift detection techniques
Data Quality and Diversity	Poor AI insights due to inconsistent or noisy telemetry	Enforce data schema standards; improve logging practices; deploy data validation layers
Explainability and Trust	Limited operator confidence in AI outputs; compliance risks	Integrate Explainable AI (XAI) modules; use interpretable models or post-hoc explanation tools
Integration with Legacy Systems	Higher engineering effort; increased time-to-value	Use adapters, service meshes, or API gateways; isolate legacy domains with custom agents
Cost and Resource Overhead	Increased compute/storage usage; budget constraints	Optimize model complexity; leverage autoscaling and serverless ML where possible
Human-in-the-Loop Operationalization	Operational friction in critical environments	Design approval workflows; integrate with policy engines; provide override capabilities

Another critical trend is the decentralization of intelligence through edge AI. With the proliferation of edge devices across sectors like industrial IoT and 5G, the need for real-time, localized inference has grown. Federated learning enables devices to collaboratively train models while keeping data local, thus enhancing privacy and reducing communication overhead [10].

Explainability and fairness remain pressing concerns. As AI systems become more autonomous, their ability to provide clear, auditable justifications for decisions is vital for gaining user trust and ensuring accountability [11]. Future AIOps platforms will likely adopt explainable AI (XAI) methods and causal reasoning frameworks to address transparency and ethical compliance.

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

ISO 3297:2007 Certified ⁜ Impact Factor 7.12 ⁜ Vol. 9, Issue 11, November 2022

DOI: 10.17148/IARJSET.2022.91122

Security in AI pipelines is another growing challenge. Adversarial attacks can exploit model vulnerabilities or manipulate data inputs, leading to compromised decision-making. Research into adversarial robustness, trusted pipelines, and zero-trust security architectures will be essential to ensure the resilience of AIOps in mission-critical environments [12].

Policy-aware orchestration is expected to gain momentum through integration with DevSecOps and automated compliance systems. This enables infrastructure to self-govern and adapt based on dynamically evaluated policies and risk models [13]. AI-powered policy enforcement engines could help maintain alignment with evolving governance, regulatory, and security requirements.

Lastly, emerging computational paradigms such as quantum computing and neuromorphic architectures may enhance the scalability and energy efficiency of future AIOps systems [14]. Although currently at a nascent stage, these technologies offer promising avenues for accelerated learning and decision-making in distributed systems.

In summary, the next generation of AIOps will likely evolve toward systems that are not only autonomous and intelligent, but also secure, interpretable, and ethically grounded. Achieving this vision will require interdisciplinary research across AI, infrastructure, cybersecurity, and human-computer interaction.

VIII. CONCLUSION

As cloud native architectures continue to scale in complexity, driven by the proliferation of microservices, containers, and distributed systems, traditional IT operations approaches have become increasingly inadequate. This shift has underscored the critical role of AI driven automation, or AIOps, as a transformative framework for ensuring system resilience, observability, and intelligent operational decision making.

This paper has explored the evolution and architectural layers of AIOps, highlighting how telemetry, machine learning, and automated orchestration collectively enable self-managing infrastructure. The proposed layered framework, alongside practical implementations from major cloud providers such as AWS, Google Cloud, and Microsoft Azure, demonstrates the real world viability of AIOps in both enterprise and open source ecosystems.

Despite these advancements, key challenges persist. Issues such as model drift, data quality degradation, limited explainability, integration with legacy environments, and operational cost overheads require ongoing attention. Addressing these limitations necessitates the development of robust system designs, continuous model retraining, and effective human oversight mechanisms.

Looking forward, research in federated learning, edge intelligence, explainable AI, and policy aware orchestration offers promising directions for expanding the capabilities of AIOps. Additionally, emerging technologies such as quantum computing and neuromorphic hardware may further enhance the performance and scalability of future infrastructure management platforms.

In conclusion, AIOps is poised to become a foundational element in the operation of modern cloud native systems. Realizing its full potential will depend on interdisciplinary collaboration across AI research, systems engineering, cybersecurity, and operational practice.

REFERENCES

[1]. Justin Garrison and Kris Nova. *Cloud Native Infrastructure*. O'Reilly Media, 2020.
[2]. John Arundel and Justin Domingus. *Cloud Native DevOps with Kubernetes*. O'Reilly Media, 2021.
[3]. Xi Cheng, Ling Zhang, and Jiayin Xu. Aiops-based anomaly detection for cloud-native environments. *Journal of Cloud Computing*, 10(1):1–15, 2021.
[4]. Gartner. Market guide for aiops platforms, 2021. https://www.gartner.com/en/documents/3989912.
[5]. Shreshth Tuli, Sourav Basu, Shikhar Tuli, and Rajkumar Buyya. Healthfog: A novel framework for health data analytics using deep learning model in fog environment. *Future Generation Computer Systems*, 104:187–200, 2020.
[6]. Amazon Web Services. Introducing amazon devops guru, 2021. https://aws.amazon.com/devops-guru/.
[7]. Google Cloud. Google cloud operations suite, 2021. https://cloud.google.com/products/operations.
[8]. Sofoklis Katsikas, Tommaso Cucinotta, and Foutse Khomh. Towards self-healing distributed systems using artificial intelligence. *Journal of Systems and Software*, 162:110516, 2020.
[9]. Tshilidzi Marwala. *Explainable AI in Industry*. Springer, 2021.

DOI: 10.17148/IARJSET.2022.91122

[10]. Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.

[11]. Andreas Holzinger, Peter Kieseberg, Edgar Weippl, and A Min Tjoa. Toward robust and trustworthy ai. *Interna- tional Journal of Information Management*, 57:102282, 2021.

[12]. Xingjun Yuan, Pei He, Qile Zhu, and Xiaolin Li. Adversarial examples: Attacks and defenses for deep learning.

[13]. *IEEE Transactions on Neural Networks and Learning Systems*, 30(9):2805–2824, 2019.

[14]. Baolei He, Jiaqi Guo, Hui He, and Hai Jin. Automated compliance checking in devsecops: Research opportunities and challenges. *ACM Computing Surveys*, 54(8):1–35, 2021.

[15]. John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

