The 2025 Colloquium for Information Systems Security Education (CISSE)

AI-Driven Cloud Security: AIOps for Threat Detection and Compliance

Authors: Advait Patel, Vaishnavi Gudur, Prashanthi Matam, Charit Upadhyay, Aparna Achanta, Swara Dave, Shalini Sudarsan

Presentation by : Shalini Sudarsan, Prashanthi Matam

Nov 12, 2025

Outline

- Motivation & Problem Statement
- Why it matters?
- AIOps
- Evolution: SIEM → SOAR → AIOps → LLMOps
- AI-Driven Cloud Security
- Challenges
- Future Direction
- Conclusion

The Challenge: Traditional Security Falls Short

The Problem

- Exploding telemetry
- Tooling mismatch
- Identity weaknesses
- Data silos
- Business Impact

Why It Matters

- Alert fatigue → missed threats
- Delayed anomaly detection
- Operational drag & cost
- Reliability at risk

Enter AIOps: A Transformative Solution

Artificial Intelligence for IT Operations (AIOps) has emerged as a transformative paradigm by integrating machine learning, predictive analytics, and automation with big data processing.

Deeper Observability

Provides comprehensive system visibility across distributed workloads

Proactive Detection

Identifies anomalies before they escalate into critical incidents

Intelligent Response

Automates incident response and initiates self-healing actions in real time

Proven Impact: The Numbers Speak

45%

Faster Detection

Reduction in mean time to detection (MTTD) compared to traditional systems

60%

Audit Efficiency

Improvement in audit preparation processes and compliance workflows

50%

Recovery Speed

Reduction in mean time to recovery with AI-powered SOCs

Empirical research demonstrates that AIOps-driven operations significantly outpace traditional monitoring systems. Organizations using AI-enhanced threat detection reduced detection latency from **24 hours to 13 hours**, marking a transformative improvement in security posture.

The Evolution of Security Operations

1

Early 2000s: SIEM Era

Security Information and Event Management systems pooled logs and performed correlation analysis based on handwritten rules and signatures. High false positives and poor predictive intelligence.

2

2010s: SOAR Systems

Security Orchestration, Automation, and Response tools automated routine workflows but remained largely reactive with pre-established rules.

3

2015-2020: AIOps Emergence

Machine learning applied to performance monitoring, anomaly detection, and predictive maintenance in cloud-native architectures.

4

2020-2023: Cloud Security Integration

AIOps became indispensable for multi-cloud and hybrid-cloud ecosystems, with vendors embedding ML directly into SIEM and SOAR offerings.

5

2023-Present: LLMOps Era

Large language models incorporated into AIOps workflows, enabling natural language threat hunting and AI-augmented security operations centers.

Comparing Security Paradigms

Paradigm	Core Capabilities	Strengths	Limitations	2024 Adoption
SIEM	Log collection, rule-based correlation, alerting	Established visibility; compliance reporting	Static rules; high false positives; blind to zero-day threats	~80% (declining)
SOAR	Workflow automation, playbooks, SIEM integration	Accelerated response; reduced manual workload	Dependent on SIEM data quality; limited adaptive detection	~45%
AIOps	ML anomaly detection; event correlation; predictive analytics	Reduces MTTD/MTTR; handles high-volume telemetry	Requires quality training data; explainability challenges	~30% (rapid growth)
LLMOps	LLM-based log parsing, natural language investigation, AI copilots	Human-like reasoning; interpretable narratives	Adversarial prompt susceptibility; high computational cost	<10% (early phase)

Four Pillars of AI-Driven Cloud Security

1. Threat Detection Layer

- Real-time processing of massive log volumes
- Temporal analysis of user and entity behavior
- Insider threat detection through access pattern analysis

2. Compliance Layer

- Continuous configuration validation
- Automated evidence gathering
- 40-60% reduction in audit preparation time

3. Response Layer

- Dynamic playbook evolution
- Self-healing automation
- Autonomous asset isolation and vulnerability patching

4. Governance Layer

- Attention visualization and SHAP values
- Counterfactual explanations
- Alignment with NIST 800-53 and ISO 27001

Advanced Anomaly Detection Techniques

Deep Learning Approaches

• Rapid shift to ML: Cloud-native telemetry accelerated anomaly-detection research beyond rules and signatures.

• Structured logs + time context: Deep models operate on log templates and learn temporal dependencies (RNNs/Transformers).

• Self-supervised Transformers: Pretraining on unlabeled logs improves generalization across systems and reduces labeling needs.

• Strong benchmark results: Reported F1 > 0.90 on HDFS, BGL, Thunderbird, and OpenStack logs—indicating fewer false positives and earlier detection.

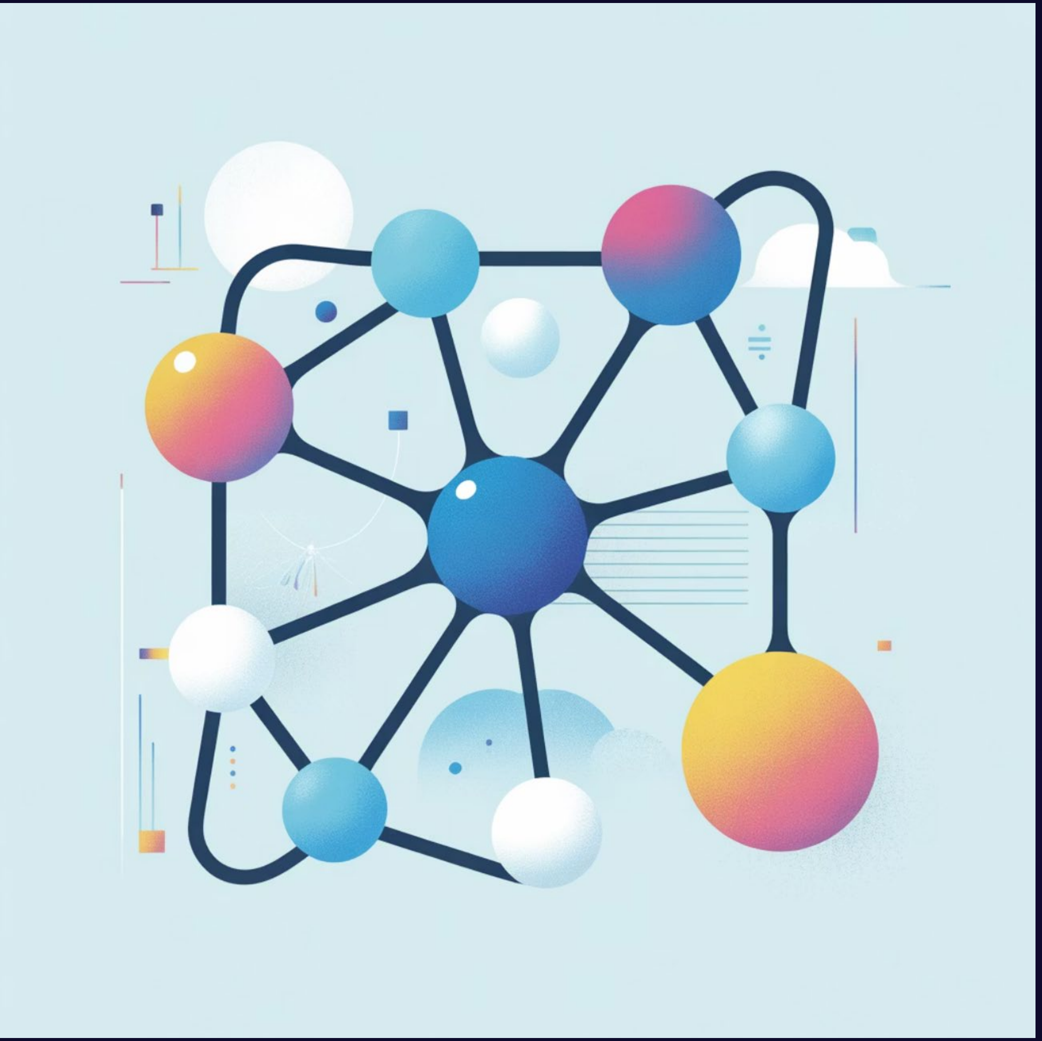

Graph-Based Methods

Graph neural networks (GNNs) construct dynamic graphs of services, containers, users, and API calls to identify anomalous subgraphs that may appear benign when viewed in isolation—such as lateral movement patterns.

LLM-Assisted Security Operations

Recent works explore Large Language Model-assisted log parsing and reasoning to convert heterogeneous, vendor-specific messages into normalized templates and enrich events with human-readable context.

Natural Language Threat Hunting

LLMs enable security analysts to query systems using natural language, dramatically reducing the learning curve and accelerating investigation workflows.

Automated Log Parsing

LLMs help generalize across unseen log formats, reducing parser maintenance and supporting few-shot pattern induction from minimal examples.

AI Copilots for SOCs

AI-augmented security operations centers where analysts interact with copilots that assist in investigation, remediation, and decision-making.

Challenge: LLMs face hallucination risk, prompt-sensitive variance, high computational costs, and privacy constraints when logs contain secrets or PII. Mitigations include RAG pipelines, prompt hardening, and policy-aware redaction.

Explainability: Building Trust in AI Security

As AIOps systems increasingly drive triage and containment actions, explainability has become central for analyst trust, governance, and post-incident forensics.

Attention Mechanisms

Attention-based architectures expose salient tokens, fields, or time steps linked to model decisions, helping analysts understand which elements contributed to risk scores.

Feature Attribution

Model-agnostic tools like SHAP and LIME provide feature importance, supporting both human validation and regulatory traceability.

Counterfactual Explanations

Show minimal changes needed to flip decisions, enabling auditors to see minimal remediation steps consistent with standards.

> "Explainability for security differs from standard XAI because analysts need **operationally actionable narratives**, not just variable importance."

Continuous Compliance Automation

The Compliance Challenge

- Manual audits are untenable in multi-cloud environments. Organizations must continuously evaluate controls against frameworks like GDPR, HIPAA, PCI DSS, ISO 27001, SOC 2, and NIST 800-53.

AI-Driven Solution

- State-of-the-art approaches encode mandates as policy-as-code, using engines such as OPA/Rego integrated into Terraform and Pulumi pipelines.

- ML components predict control failure risk by learning from historical misconfigurations, deployment patterns, and ticket backlogs.

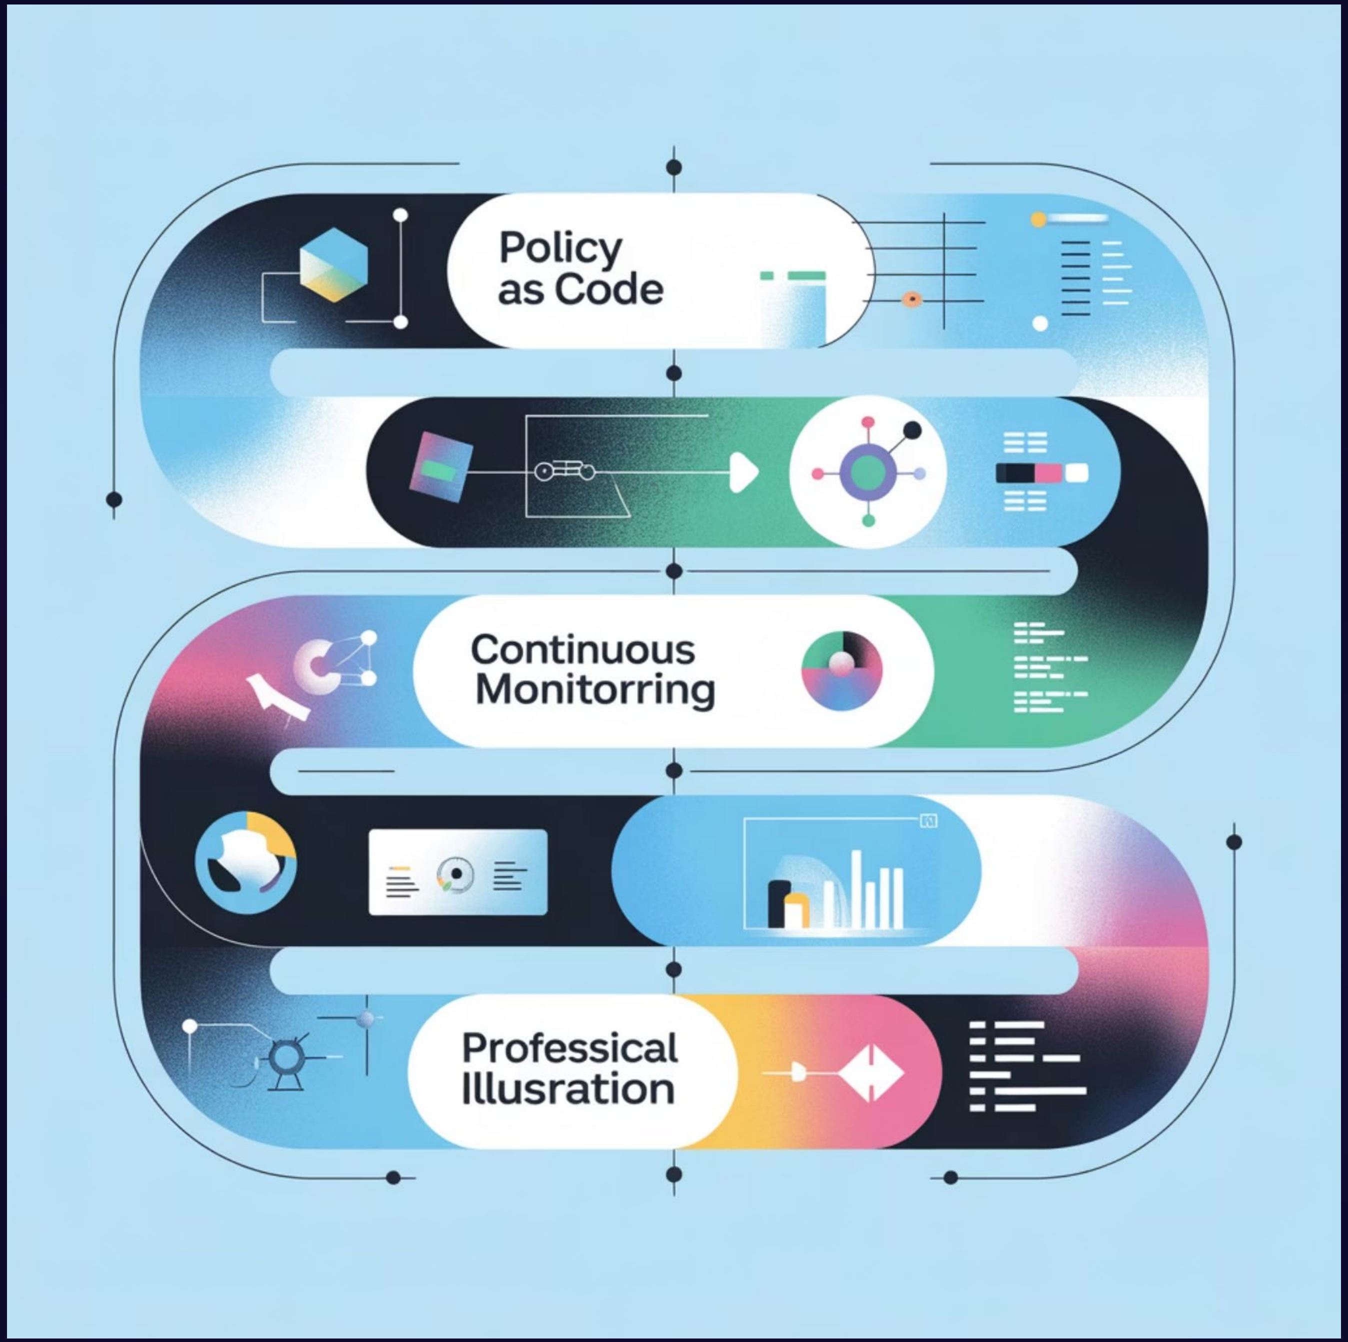

Adoption Trends and Market Growth

- MTTD reduced from ~24h to ~13h.
- Audit prep reduced by 40–60% via continuous evidence.
- Rapid enterprise adoption and vendor integration (2022–2026 trend).

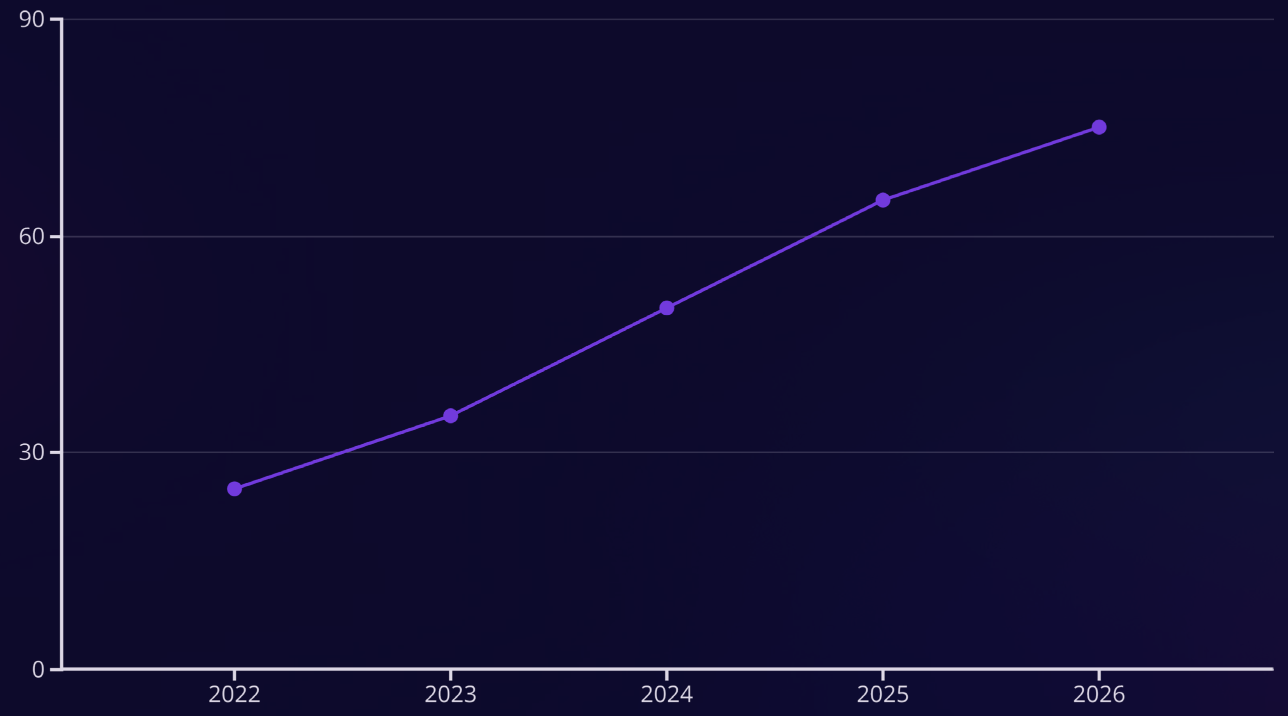

The global AIOps market is expanding at a compound annual growth rate exceeding **25 percent** from 2023 to 2030, reaching an estimated valuation of **$65 billion USD** by the end of the decade.

Critical Challenges Ahead

Adversarial Vulnerability

Attackers may deliberately poison training data or craft adversarial inputs that evade detection, undermining the trustworthiness of AI-enhanced SOCs and demanding robust defenses against model manipulation.

Explainability Gap

Security analysts struggle to interpret opaque AI-driven alerts. Absence of transparency can cause incident responders to doubt systems even in high-stakes situations.

Algorithmic Bias

AI models may disproportionately classify certain anomaly types, resulting in distorted operational responses and higher rates of false positives for specific user groups or regions.

Multi-Cloud Complexity

Integration of AIOps pipelines in multi-cloud and multi-vendor settings presents interoperability issues, where security platforms vary in data structures, APIs, and policy frameworks.

Regulatory Fragmentation

The fractured landscape of compliance systems such as GDPR, HIPAA, and PCI DSS poses challenges to systematic approaches for automated compliance across various jurisdictions.

Future Directions: The Road Ahead

LLM-Based SOC Agents

Autonomous research and triage through Retrieval-Augmented Generation (RAG) and real-time log information, serving as digital co-pilots to human analysts.

AI Governance Platforms

Implementation of fairness, explainability, and compliance systems at the AIOps pipeline level, ensuring accountable and auditable automation.

Vendor-Agnostic Unified Platforms

Cross-cloud threat intelligence and uniform security postures regardless of provider policies, addressing the multi-cloud challenge.

Self-Healing Infrastructures

AI systems that not only detect and react to incidents but actively repair vulnerabilities and restore compliance—the end-stage of AIOps-driven defense.

These paths underscore the convergence of AI and cloud-native systems, pointing toward resilient, adaptive cybersecurity environments capable of autonomous defense and continuous improvement.

Conclusion: A New Era of Cloud Security

AIOps shifts security – Reactive to predictive.

Meaningful gains in speed, scale, and compliance readiness.

Trustworthiness (XAI, fairness, robustness) is the gating factor

Action: start with detection + compliance-as-code; stage automation

Thank You

Questions & Discussion