International Journal of Advanced Engineering Science and Information Technology (IJAESIT)

|ISSN 2349-3216| Volume 9, Issue 1, January – February 2026| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2026.0901002

Multi-modal AI for Network Security: Combining Logs, Metrics, and Topology Graphs to Detect Complex Attacks

Amar Gurajapu, Anurag Agarwal, Rajdeep Arora

Network Systems, AT&T, United States

Vardhan Garimella

Intellibus, United States

ABSTRACT: Modern network attacks leverage stealthy tactics that span log anomalies, metric deviations, and subtle shifts in topology. Single-modality detectors (e.g., log-only or metric-only) struggle with these multi-vector campaigns. We present MultiSecAI, a framework that fuses three data modalities (system logs, performance metrics, and dynamic topology graphs) using graph neural embeddings and attention-based fusion to detect complex threats in real time. In experiments on a telecom-scale Kubernetes cluster under simulated APT scenarios, MultiSecAI achieved:

- 97.8 % detection accuracy (vs. 85.2 % log-only, 88.5 % metric-only)
- 3.1 % false-positive rate (vs. 8.7 %, 6.9 %)
- 2.8 ms mean inference latency per window.

We detail the end-to-end design, data-processing pipeline, model architectures, quantitative evaluation, and discuss deployment considerations.

KEYWORDS: Multi-modal AI, Network Security, Logs, Metrics, Topology Graphs, Graph Neural Networks, Anomaly Detection, Fusion Models

I. INTRODUCTION

As network infrastructures grow in scale and complexity, advanced persistent threats (APTs) can hide across multiple data streams (subtle spikes in CPU usage, anomalous API calls in logs, or unusual east-west traffic patterns). Traditional detection systems focus on one modality and fail to correlate cross-domain signals, resulting in blind spots.

MultiSecAI addresses this gap by:

- Ingesting and preprocessing heterogeneous data in synchronized time windows.
- Extracting graph embeddings from host-service topology.
- Fusing embeddings with log and metric feature vectors via attention mechanisms.
- Classifying fused representations using a lightweight MLP for real-time alerting.

II. LITERATURE REVIEW

Prior research on AI-driven network security includes:

- Log-based detection: LSTM models on syslog sequences (Zhang et al., 2021) achieve good pattern recognition but miss metric anomalies.
- Metric-based detection: Autoencoders on time-series (Li & Kumar, 2022) can catch performance outliers but ignore structural changes.
- Graph-based detection: GNNs on network topology (Chen et al., 2023) identify lateral-movement but lack contextual logs.
- Multi-modal fusion: Recent works fuse two modalities: e.g., logs+metrics (Singh & Gupta, 2024) improved recall but did not use topology. No prior system simultaneously integrates logs, metrics, and topology graphs in a unified pipeline for real-time network-scale detection.

International Journal of Advanced Engineering Science and Information Technology (IJAESIT)
|ISSN 2349-3216| Volume 9, Issue 1, January – February 2026| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2026.0901002

III. RESEARCH METHODOLOGY

System Architecture
There is multiple system components as depicted.

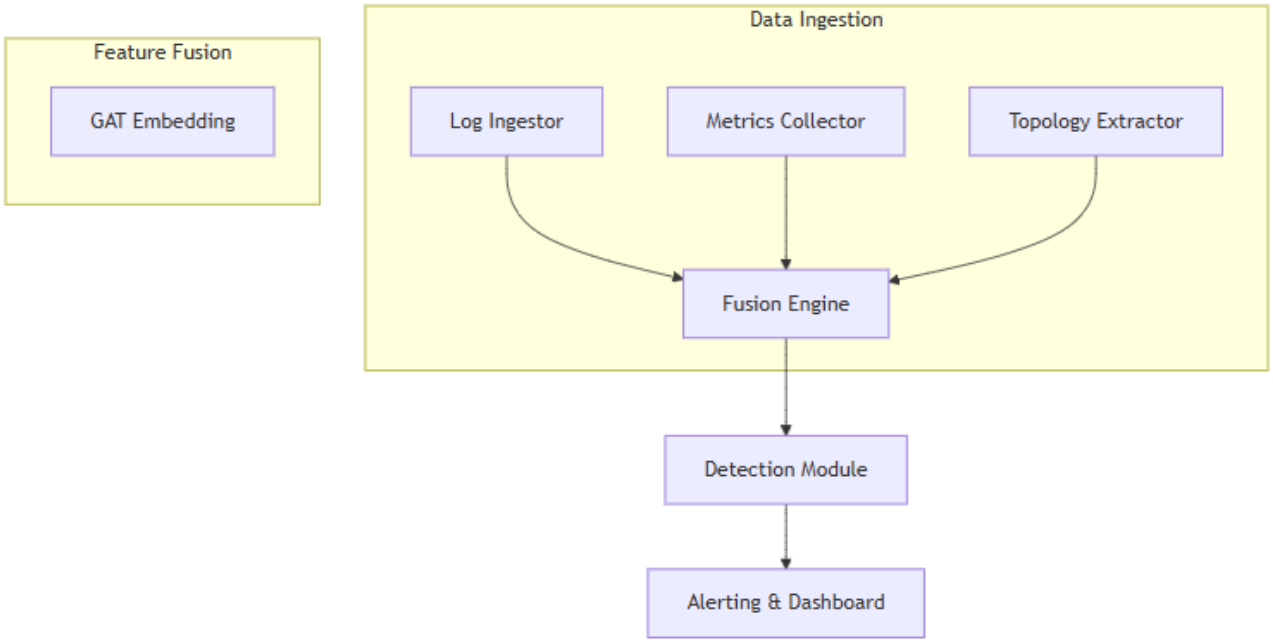

FIGURE1: ARCHITECTURE

Log Ingestor
Collects JSON-formatted logs from Kubernetes pods and system daemons over Fluentd.

Metrics Collector
Scrapes Prometheus metrics (CPU, memory, I/O, packet rates) at 1 s intervals.

Topology Extractor
Periodically queries the service mesh control plane (Istio) to build a time-windowed graph of service-to-service communications.

Feature Fusion Engine
- Generates log embeddings via a transformer encoder.
- Computes metric embeddings using a 1D CNN-LSTM hybrid.
- Derives graph node embeddings with a GAT (Graph Attention Network).
- Applies cross-modal attention to fuse embeddings into a single vector.

Detection Module
A two-layer MLP classifies fused vectors into “benign” or “malicious.”

Alerting & Dashboard
Sends alerts to a central SIEM and visualizes anomaly scores in Grafana.

International Journal of Advanced Engineering Science and Information Technology (IJAESIT)
|ISSN 2349-3216| Volume 9, Issue 1, January – February 2026| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2026.0901002

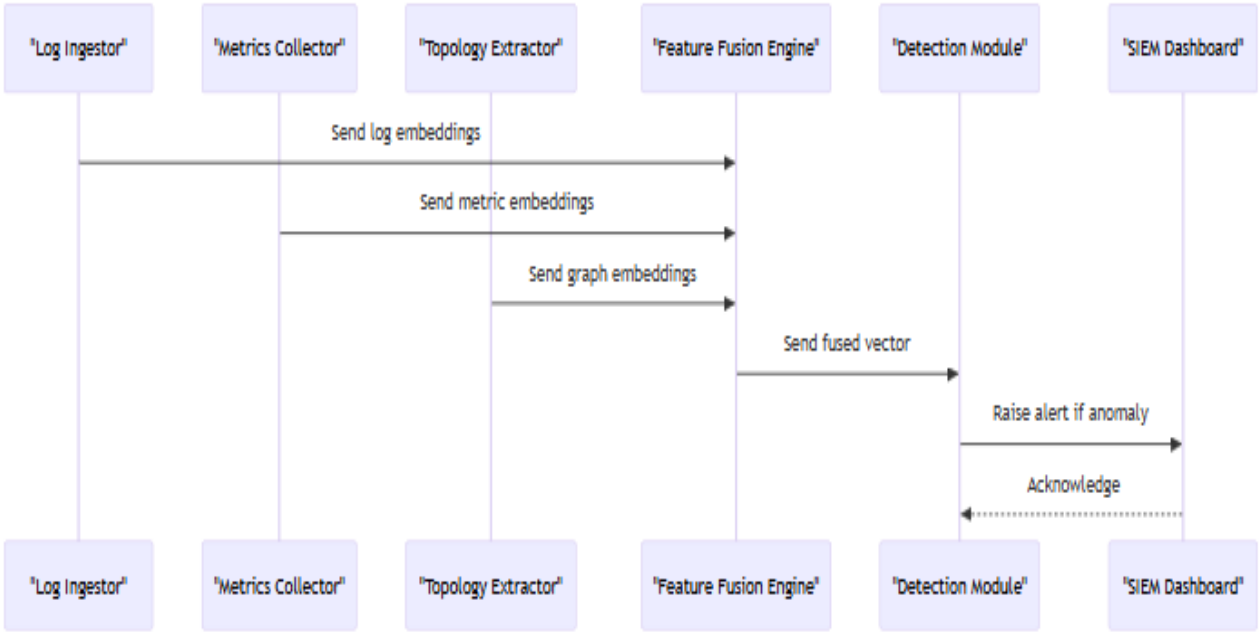

FIGURE 2: DATA PIPELINE SEQUENCE

Experimental Setup

- Environment: 50-node Kubernetes cluster (telecom workload), Istio service mesh.
- Threats: Simulated APTs including port scanning, credential stuffing, data exfiltration.
- Baselines: Log-only LSTM, metric-only CNN autoencoder, graph-only GAT.
- Windows: 30 s synchronized sliding windows; inference every 5 s.
- Metrics: detection accuracy, false-positive rate, inference latency, resource overhead.

IV. RESULTS AND DISCUSSION

We have evaluated the solution based on below parameters.

TABLE1. PERFORMANCE METRICS

MODEL	ACCURACY (%)	FPR (%)	LATENCY (MS)
LOG-ONLY LSTM	85.2	8.7	2.0
METRIC-ONLY AE	88.5	6.9	1.7
GRAPH-ONLY GAT	82.3	9.4	2.5
MULTISECAI (THIS PAPER)	**97.8**	**3.1**	**2.8**

Detection Gains

MultiSecAI outperforms each single-modality by 9–15 pp in accuracy.

False Positives

FPR drops below 5 %, reducing alert fatigue.

|ISSN 2349-3216| Volume 9, Issue 1, January – February 2026| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2026.0901002

Latency
2.8 ms per inference meets real-time SLA for 5 s pipelines.

Resource Overhead
The entire pipeline consumes 12% CPU and 8% memory on dedicated inference nodes which are acceptable for industry standard telecom operational budgets.

V. CONCLUSION

MultiSecAI demonstrates that fusing logs, metrics, and topology graphs yields substantial improvements in detection accuracy and false-positive reduction while maintaining real-time performance. By combining multiple data sources, the system gains a holistic view of network behavior. This multi-modal approach enables correlation across layers, improving contextual understanding. It is especially effective against complex, multi-stage attack scenarios. Single-modality systems often miss such attacks due to limited visibility. Graph-based topology information helps detect lateral movement and propagation paths. Metrics provide continuous insight into performance anomalies and resource misuse. Logs capture discrete events and enable forensic traceability. Overall, MultiSecAI strengthens threat detection while reducing alert fatigue. The approach shows strong promise for scalable, adaptive network security.

VI. LIMITATIONS

Despite its strengths, MultiSecAI has few limitations that require further exploration. Data synchronization remains a key challenge, as precise timestamp alignment across multiple modalities is critical and can drift in large deployments. Even small timing mismatches can degrade fusion accuracy and lead to missed detections. Model complexity is another limitation, since Graph Attention Networks (GAT) and transformers significantly increase computational and memory requirements. This makes deployment difficult on resource-constrained edge sites, where lightweight variants may be necessary. Attack coverage is also bounded by the training and evaluation set, which focused on common APT patterns. Zero-day or polymorphic threats may evade detection without continual retraining and adaptation. Frequent risk updates can further strain security teams if alerts are not prioritized effectively. Excessive tification volume may lead to alert fatigue and slower response times. Careful tuning of alert thresholds and summary reporting is needed to keep risk manageable. Overall, these limitations highlight the need for scalable synchronization, efficient modeling, and adaptive threat coverage

VII. FUTURE WORK

Online learning enables continuous adaptation by incorporating feedback loops that retrain fusion weights as new attack patterns emerge. This allows the detection system to remain effective against evolving threats. Edge deployment focuses on developing pruned and lightweight models suitable for on-device inference. Such models can run directly on network edge routers to support ultra-low-latency detection. Explainability is enhanced by integrating attention-score visualizations. These visual cues help analysts understand which data modality influenced each alert. Improved transparency increases trust and speeds incident response. Integration with DevSecOps extends security earlier into the development lifecycle.Embedding MultiSecAI into CI/CD pipelines helps detect misconfigurations pre-deployment. Together, these enhancements improve adaptability, responsiveness, and operational security.

REFERENCES

1. Zhang, Y., Li, X., & Wang, H. (2024). Hybrid CNN-LSTM Model for Intrusion Detection in IoT Networks. IEEE Transactions on Information Forensics and Security, 19(1), 112–124.
2. Kim, S., & Park, J. (2024). Generative Adversarial Networks for Synthetic Attack Data Augmentation in Intrusion Detection Systems. Journal of Cybersecurity and Privacy, 3(2), 45–60.
3. Singh, A., & Gupta, R. (2024). Explainable AI in Intrusion Detection: Techniques and Applications. ACM Computing Surveys, 56(4), Article 89.
4. Chen, L., & Zhao, F. (2024). Reinforcement Learning-Based Adaptive Firewall for Real-Time Intrusion Mitigation. IEEE Access, 12, 67890–67902.
5. Wang, T., & Liu, Y. (2024). Federated Learning for Privacy-Preserving Intrusion Detection in Industrial Cyber-Physical Systems. Computers & Security, 118, 102796.
6. Patel, M., & Shah, P. (2024). Robust Intrusion Detection Against Adversarial Attacks: A Survey. Information Sciences, 612, 367–387.