| Volume 02 | Issue 06 | November 2026 | ISSN: 3139-1737

International Journal of Advances and Applications in Convergent Studies & Transdisciplinary (IJ-AACST)

Bridging Disciplines, Bridging Knowledge

Aligned with the Vision of NEP 2020 and UGC Guidelines

Website: https://aacstij.org || Email: journaleditor@aacstij.org

Sentinel LLM: An Evidence-Grounded Hybrid AI Framework for Cloud Threat Detection and Risk-Aware Incident Response

DOI: https://doi.org/ 10.5281/zenodo.24103862

S. D. Gujar[1]*, Dr. V. S. Tondre[2]

*Corresponding Author: [1]Department of Computer Science & Application, Brijlal Biyani Science College, Amravati, Maharashtra, India; affiliated to Sant Gadge Baba Amravati University (SGBAU), Amravati

santooshgujar@gmail.com

ORCID iD: https://orcid.org/0009-0006-3203-6747

[2]Department of Computer Science & Application, Brijlal Biyani Science College, Amravati, Maharashtra, India; affiliated to Sant Gadge Baba Amravati University (SGBAU), Amravati

vstondre@gmail.com

ORCID iD: https://orcid.org/0009-0001-1357-4527

ABSTRACT

Cloud platforms produce large volumes of security information, including login records, API calls, network events, permission changes, and file-access logs. Rule-based systems may miss unusual behaviour or create too many alerts. Large language models (LLMs) can explain technical events in simple language, but they may also produce unsupported statements. This paper proposes SentinelLLM, a hybrid framework that combines machine-learning anomaly detection, retrieval-augmented generation (RAG), and an LLM for cloud threat analysis. The framework collects cloud events, prepares the data, identifies unusual activity, retrieves relevant security policies, and generates an evidence-linked explanation and response recommendation. High-impact actions require policy checks and human approval. An AWS prototype is proposed using CloudTrail, GuardDuty, CloudWatch, S3, Lambda, SQS, and Amazon Bedrock. The paper defines an evaluation plan based on detection quality, evidence support, response quality, latency, and cost. Numerical examples in the result tables are explicitly illustrative planning values; they are not measured experimental results.

KEYWORDS: Cloud security, large language models, anomaly detection, retrieval-augmented generation, AWS, threat detection, incident response, explainable artificial intelligence.

1. INTRODUCTION

Cloud computing is widely used by businesses, colleges, hospitals, government offices, and online services. It provides computing power, storage, and applications without requiring every organization to maintain its own hardware. However, cloud environments create security challenges because resources, users, and permissions can

© 2026 The Author(s). This article is published by the AACST in the IJ-AACST under the Creative Commons Attribution 4.0 International (CC BY 4.0) License. The authors retain copyright and grant AACST a non-exclusive license to publish, distribute, and archive the work.