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

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0706010

Large Language Model-Based Intelligent Cyber Threat Analysis for Sovereign Cloud Infrastructure

Mark Richards

Software Architect, Independent Consultant, London, United Kingdom

ABSTRACT: Sovereign cloud infrastructure is becoming increasingly important for governments, regulated organizations, critical industries, and enterprises seeking greater control over sensitive data, infrastructure, digital services, and cybersecurity operations. However, sovereign cloud environments face increasingly complex cyber threats involving advanced persistent attacks, identity compromise, malicious insiders, supply-chain vulnerabilities, API exploitation, ransomware, and sophisticated lateral movement. Large Language Models (LLMs) provide new opportunities for intelligent cyber threat analysis because they can process heterogeneous security information, understand contextual relationships among events, summarize incidents, generate threat hypotheses, and assist analysts in security decision-making. This paper proposes an LLM-based intelligent cyber threat analysis framework specifically designed for sovereign cloud infrastructure. The framework integrates security logs, network telemetry, identity events, vulnerability information, cloud configuration data, threat intelligence, and incident reports into an AI-driven analytical pipeline. Retrieval-Augmented Generation (RAG), domain-specific knowledge bases, structured threat representations, and explainable reasoning mechanisms are used to improve analytical accuracy and reduce unsupported outputs. A sovereignty-aware governance layer controls data residency, model access, auditability, privacy, and policy enforcement. The proposed methodology evaluates threat detection accuracy, contextual relevance, false-positive reduction, response latency, explanation quality, and operational overhead. The framework aims to strengthen sovereign cloud security by providing intelligent, auditable, privacy-aware, and context-sensitive cyber threat analysis while maintaining organizational control over sensitive security information and AI operations.

KEYWORDS: Large Language Models, Cyber Threat Analysis, Sovereign Cloud, Cloud Security, Artificial Intelligence, Threat Intelligence, Retrieval-Augmented Generation, Cybersecurity, Data Sovereignty, Zero Trust, Security Operations, Privacy-Preserving AI

I. INTRODUCTION

The rapid expansion of cloud computing has transformed digital infrastructure by enabling organizations to dynamically provision computing resources, deploy applications at scale, and access advanced data-processing capabilities. However, the growing dependence on cloud infrastructure has also introduced concerns surrounding data sovereignty, regulatory compliance, national security, operational independence, and control over critical digital resources. Sovereign cloud infrastructure addresses these concerns by emphasizing jurisdictional control, data residency, trusted infrastructure operations, regulatory alignment, and organizational authority over sensitive workloads and information. Governments, financial institutions, healthcare organizations, defense-related industries, and critical infrastructure providers increasingly require cloud environments where sensitive information remains within specified geographic and legal boundaries. Although sovereign cloud architectures can improve governance and control, they remain exposed to complex cybersecurity threats. Attackers may target cloud management interfaces, identity systems, application programming interfaces, containers, virtual machines, storage systems, supply chains, and administrative accounts. Security teams must therefore analyze enormous quantities of heterogeneous information generated across infrastructure, applications, networks, identities, and cloud services. Conventional security information and event management systems primarily depend on predefined rules, signatures, correlation logic, and statistical thresholds. These mechanisms remain valuable but can struggle to understand complex relationships among distributed security events. Large Language Models offer an emerging capability for interpreting natural-language and structured security information, correlating related observations, summarizing incidents, and supporting threat investigation. An LLM-based cyber threat analysis architecture can potentially transform fragmented security telemetry into contextual intelligence while maintaining sovereignty-oriented controls over sensitive information.

Cyberattacks against sovereign cloud infrastructure can be particularly consequential because compromise may affect government services, national critical infrastructure, regulated datasets, or strategic enterprise operations. Attackers may employ multi-stage techniques in which individual events appear relatively harmless but become suspicious when

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0706010

considered collectively. For example, an unusual authentication event followed by privilege escalation, access to an uncommon API, abnormal data retrieval, and communication with an unfamiliar external endpoint may represent a coordinated intrusion. Detecting such campaigns requires understanding temporal and semantic relationships across different sources. LLMs are particularly relevant because they can process diverse textual information such as security alerts, audit logs, incident reports, threat intelligence descriptions, vulnerability documentation, configuration records, and analyst notes. Instead of simply identifying an individual anomaly, an LLM can assist in constructing an incident narrative and connecting observations to known attack behaviors. Retrieval-Augmented Generation can further improve reliability by retrieving information from trusted organizational knowledge bases before generating analytical responses. This can include security policies, cloud architecture documentation, vulnerability databases, incident-response procedures, asset inventories, and approved threat-intelligence sources. However, applying general-purpose LLMs directly to sovereign security data introduces significant concerns. Sensitive telemetry could be exposed to external model providers, generated outputs could contain inaccurate conclusions, and model prompts may unintentionally reveal confidential information. Therefore, sovereignty-aware LLM deployment requires careful consideration of data residency, model hosting, access control, encryption, auditability, privacy, and governance.

The proposed intelligent cyber threat analysis framework addresses these challenges by treating the LLM as a controlled analytical component rather than an unrestricted autonomous security authority. The framework consists of several interconnected layers, beginning with secure telemetry collection from sovereign cloud infrastructure. Data sources include cloud audit logs, identity and access-management events, network flows, endpoint activity, application logs, vulnerability information, configuration states, container activity, API requests, and security alerts. A normalization layer converts these heterogeneous sources into structured security events. An intelligence-processing layer performs filtering, deduplication, enrichment, and contextual correlation before relevant information is passed to the LLM. A sovereign retrieval layer provides access to approved internal knowledge sources without unnecessarily transferring sensitive information outside the controlled environment. The LLM then performs tasks such as incident summarization, threat classification, attack-chain reconstruction, hypothesis generation, root-cause analysis, and analyst question answering. A policy and governance layer evaluates model outputs against organizational rules and prevents unauthorized disclosure or unsafe recommendations. Explainability mechanisms record the evidence used to support analytical conclusions. The framework can also integrate Zero Trust principles by continuously evaluating identities, workloads, access requests, and communication patterns. This architecture allows organizations to benefit from LLM capabilities while retaining control over sensitive information and security operations.

Another important consideration is the relationship between intelligent analysis and sovereign operational independence. A sovereign cloud should not become dependent on external AI infrastructure for critical security functions if such dependence creates unacceptable jurisdictional, availability, or confidentiality risks. Therefore, the proposed framework emphasizes deployment flexibility, including private LLM instances, locally hosted models, controlled inference environments, and sovereign cloud-based AI services. Model selection should consider not only language performance but also security, transparency, computational requirements, update mechanisms, and governance requirements. Fine-tuning can adapt models to domain-specific cybersecurity terminology, while RAG can provide current information without requiring constant model retraining. LLM outputs should be treated as analytical recommendations whose reliability is measured using evidence-grounded evaluation. For high-impact actions, human analysts should remain responsible for final authorization. Lower-risk activities such as alert summarization, event clustering, evidence extraction, and investigation assistance can be more highly automated. The research therefore aims to establish an intelligent cyber threat analysis model that combines LLM reasoning capabilities with sovereign security controls, trusted information retrieval, Zero Trust principles, and human-centered governance. The ultimate objective is to improve the speed, contextual quality, and consistency of cybersecurity analysis while ensuring that sovereign cloud environments retain control over sensitive data, AI operations, and security decisions.

II. LITERATURE REVIEW

The development of sovereign cloud computing has been driven by increasing concerns about data sovereignty, jurisdictional exposure, regulatory requirements, and control over critical digital infrastructure. Conventional public cloud models can create challenges when sensitive information is processed or stored across multiple jurisdictions or when cloud operations depend heavily on foreign infrastructure providers. Sovereign cloud architectures seek to address these concerns through mechanisms such as geographic data residency, controlled administrative access, localized infrastructure operations, encryption, trusted execution environments, compliance controls, and organizational governance. Research in cloud security has traditionally emphasized confidentiality, integrity, availability, identity management, access control, encryption, secure configuration, and continuous monitoring. However, sovereignty introduces additional dimensions involving legal jurisdiction, operational autonomy, supply-chain trust, and control over technology dependencies. Security monitoring within sovereign environments therefore requires mechanisms

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0706010

capable of processing information locally and maintaining strict boundaries around sensitive telemetry. Existing Security Information and Event Management and Security Operations Center technologies provide centralized mechanisms for event collection and correlation, but large-scale environments can generate enormous alert volumes. This has encouraged research into artificial intelligence and machine learning for automated anomaly detection, event classification, threat prioritization, and security analytics. The increasing sophistication of cloud-native threats indicates that purely rule-based mechanisms may be insufficient for interpreting complex multi-stage attack campaigns.

Artificial intelligence-based cybersecurity research has explored supervised learning, unsupervised anomaly detection, deep learning, graph-based analysis, and behavioral analytics. Machine learning algorithms can classify network traffic, detect malicious behavior, identify suspicious authentication patterns, and prioritize security alerts. Deep learning models can identify nonlinear relationships within high-dimensional telemetry, while graph-based methods can represent relationships among users, workloads, services, identities, and resources. These approaches have improved automated detection, but they often require carefully structured datasets and domain-specific feature engineering. Large Language Models introduce a different analytical paradigm because they can interpret both structured and unstructured information using natural-language representations. Modern LLMs can perform summarization, classification, information extraction, question answering, reasoning assistance, and report generation. In cybersecurity, these capabilities have been investigated for security alert summarization, threat-intelligence analysis, malware explanation, vulnerability analysis, incident-response assistance, and security operations automation. LLMs can potentially reduce analyst workload by converting thousands of technical events into concise incident narratives. They can also assist analysts in querying security information using natural language. Nevertheless, cybersecurity applications of LLMs face significant problems involving hallucination, incomplete reasoning, prompt injection, outdated knowledge, data leakage, and susceptibility to manipulated input. These challenges are particularly serious in sovereign environments because incorrect or unauthorized AI outputs may affect critical infrastructure or expose highly sensitive information.

Retrieval-Augmented Generation has emerged as an important technique for improving the reliability and contextual relevance of LLM applications. Instead of depending entirely on information encoded during model training, RAG retrieves relevant documents from trusted knowledge repositories and provides them as contextual evidence during inference. For cybersecurity, this approach can connect LLMs with vulnerability databases, security policies, asset inventories, incident histories, threat-intelligence repositories, configuration standards, and organizational response procedures. RAG is particularly valuable in sovereign cloud environments because knowledge repositories can remain inside controlled infrastructure. The LLM can therefore analyze sensitive information without necessarily requiring external access to the organization's data. Research has also explored specialized cybersecurity knowledge graphs that represent relationships among vulnerabilities, assets, attack techniques, indicators, identities, and incidents. Combining knowledge graphs with LLMs can improve contextual reasoning by providing structured relationships that complement natural-language analysis. However, retrieval quality remains a critical factor. If irrelevant, outdated, or malicious documents are retrieved, the LLM may produce misleading conclusions. Security-focused RAG systems therefore require document provenance, access control, content validation, freshness management, and retrieval auditing. Prompt injection within retrieved documents is another concern because attackers may attempt to influence model behavior through manipulated security logs or external content.

The literature on trustworthy and secure AI highlights the importance of explainability, privacy, governance, robustness, and human oversight. Explainable AI techniques seek to provide interpretable evidence for model decisions, which is particularly important for security analysts who need to validate automated conclusions before taking disruptive actions. Privacy-preserving machine learning techniques, including data minimization, encryption, federated learning, differential privacy, and secure execution environments, have also been investigated for protecting sensitive information. For LLMs, additional controls are required because prompts and generated responses can unintentionally expose confidential information. Access-control mechanisms, prompt filtering, output validation, sensitive-data masking, audit logging, and controlled model execution can reduce these risks. Research on agentic AI further explores systems capable of planning and executing multi-step tasks, creating opportunities for autonomous security investigation and response. However, excessive autonomy can introduce significant operational risks, particularly when an AI system is allowed to modify infrastructure or revoke access without human approval. Existing research therefore supports bounded autonomy in which AI systems can recommend or execute predefined low-risk actions while high-impact decisions remain subject to policy and human authorization. These findings indicate a research gap in integrating LLM-based analysis, sovereign data governance, trusted retrieval, cloud security telemetry, and controlled response into a unified framework. The proposed research addresses this gap by developing a sovereignty-aware LLM architecture designed specifically for intelligent cyber threat analysis.

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0706010

III. RESEARCH METHODOLOGY

The proposed research methodology adopts a multi-layer experimental design for developing and evaluating an LLM-based intelligent cyber threat analysis framework for sovereign cloud infrastructure. The first stage establishes a controlled sovereign cloud test environment representing government, financial, healthcare, or critical enterprise workloads. The environment contains virtual machines, containerized services, Kubernetes workloads, databases, APIs, identity-management services, storage systems, and network components. Security telemetry is collected from cloud audit mechanisms, identity providers, firewalls, network monitoring systems, application logs, operating-system events, container runtimes, vulnerability scanners, and configuration-management systems. To preserve sovereignty, the experimental architecture keeps sensitive telemetry inside a controlled infrastructure boundary. Data is normalized into a common security-event representation containing timestamps, identities, assets, actions, network information, severity, source, and contextual metadata. Data preprocessing removes duplicate events, handles missing values, masks unnecessary sensitive information, and categorizes events according to security relevance. A threat-intelligence integration layer enriches observations using approved internal and external intelligence sources where legally and operationally permitted. The resulting dataset contains both benign operational activity and controlled attack scenarios, enabling systematic evaluation of LLM-supported threat analysis.

The second stage develops the LLM-based analytical layer using a combination of retrieval, structured context construction, and domain adaptation. Instead of sending raw security telemetry directly to the model, a preprocessing engine selects relevant events and transforms them into concise security contexts. A sovereign knowledge repository is constructed containing approved security policies, architecture documentation, vulnerability information, threat-intelligence records, asset classifications, incident-response procedures, attack-pattern descriptions, and historical security cases. Retrieval-Augmented Generation is used to retrieve relevant evidence for each investigation request. Retrieval ranking considers semantic similarity, asset relevance, event timestamps, document authority, and information freshness. The LLM receives the retrieved evidence together with structured event information and a predefined analytical prompt. It performs tasks including threat classification, incident summarization, attack-chain reconstruction, suspicious behavior explanation, risk assessment, and investigation-question generation. To reduce hallucination, generated conclusions are required to identify supporting evidence and distinguish observed facts from analytical hypotheses. Domain-specific adaptation can be performed using curated cybersecurity examples while maintaining strict controls over training-data confidentiality. Multiple LLM configurations can be evaluated, including general-purpose models, cybersecurity-adapted models, and locally deployed sovereign models.

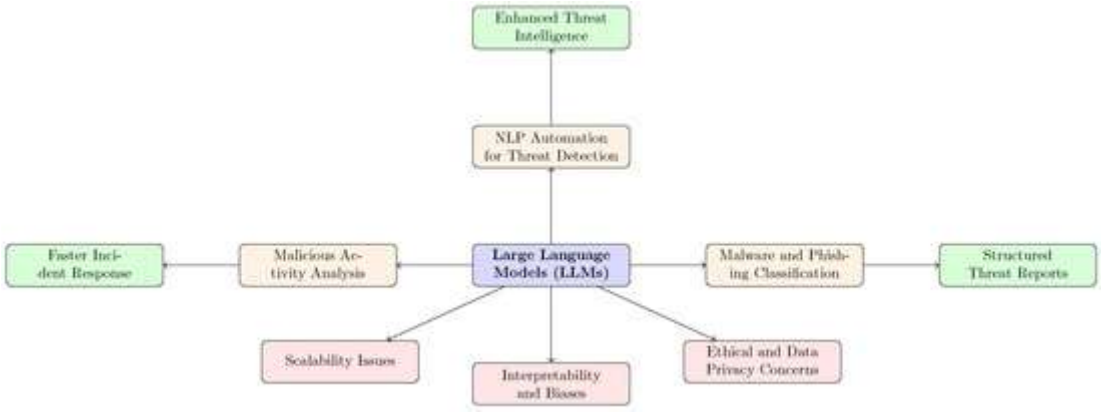


FIG1: Large Language Model-Based Intelligent Cyber Threat Analysis

The third stage develops the threat reasoning and governance mechanism. Each detected incident is assigned a contextual risk score based on factors including anomaly severity, affected asset criticality, identity privilege, vulnerability exposure, confidence level, attack-stage indicators, and potential operational impact. The LLM's interpretation is combined with deterministic security rules and structured threat intelligence rather than being treated as the sole source of truth. A security reasoning engine correlates individual observations into possible attack chains. For example, unusual authentication, privilege escalation, abnormal API usage, suspicious data access, and external communication can be analyzed as related stages of a potential intrusion. The governance layer validates model inputs and outputs, checks access permissions, detects sensitive information, and records model interactions for audit purposes. Prompt-injection defenses are implemented to prevent malicious content in logs or retrieved documents from overriding security instructions. Output validation checks whether generated recommendations are supported by available evidence. For high-risk decisions, the system generates recommendations for security analysts rather than

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0706010

executing them automatically. For low-risk activities such as alert summarization, evidence classification, and investigation assistance, greater automation can be permitted. This controlled-autonomy model preserves human accountability while benefiting from rapid AI-assisted analysis.

The fourth stage evaluates the framework through controlled cybersecurity experiments and comparative analysis. Attack scenarios include credential compromise, privilege escalation, API abuse, ransomware-like activity, unauthorized data access, cloud misconfiguration exploitation, malicious container deployment, lateral movement, supply-chain compromise, and anomalous administrative behavior. Each scenario generates a sequence of security events that must be analyzed by the proposed system. The framework is compared against conventional rule-based detection and traditional machine-learning-based security analytics. Evaluation metrics include threat classification accuracy, precision, recall, F1-score, false-positive rate, detection latency, incident summarization quality, evidence-grounding accuracy, analyst workload reduction, and response recommendation accuracy. LLM-specific evaluation also measures hallucination frequency, unsupported claims, retrieval relevance, explanation completeness, and resistance to prompt injection. Sovereignty-related measurements include data residency compliance, percentage of sensitive data retained within the controlled environment, audit completeness, model-access traceability, and dependency on external AI services. Performance and scalability are evaluated by increasing the number of security events, monitored assets, concurrent investigations, and knowledge-base documents. Ablation experiments remove individual components such as RAG, threat intelligence, contextual risk scoring, and governance controls to determine their contribution. The final analysis assesses whether the proposed framework can improve threat-analysis quality and operational efficiency while preserving confidentiality, sovereignty, auditability, and human oversight.

Advantages

1. **Context-Aware Threat Analysis:** LLMs can correlate security events and generate broader incident context rather than treating alerts independently.
2. **Intelligent Alert Summarization:** Large volumes of security alerts can be converted into concise and understandable incident summaries.
3. **Natural-Language Investigation:** Security analysts can interact with complex security telemetry using natural-language questions.
4. **Improved Threat Correlation:** The framework can connect identity, network, application, vulnerability, and cloud events to identify potential attack chains.
5. **Sovereign Data Control:** Local or controlled model deployment can help organizations retain sensitive security information within approved jurisdictions.
6. **Retrieval-Augmented Analysis:** RAG can connect LLM reasoning with current organizational policies, threat intelligence, asset information, and security documentation.
7. **Reduced Analyst Workload:** Repetitive analysis, evidence extraction, and report generation can be automated.
8. **Faster Incident Investigation:** AI-assisted reasoning can reduce the time required to understand complex security incidents.
9. **Explainable Security Analysis:** Evidence-grounded responses can help analysts understand why an incident was classified as suspicious.
10. **Integration with Threat Intelligence:** External and internal intelligence can enrich security-event interpretation.
11. **Support for Zero Trust:** LLM-based analysis can assist with understanding identity, access, workload, and communication anomalies.
12. **Scalable Security Operations:** AI can assist in analyzing high volumes of heterogeneous security information.
13. **Improved Incident Documentation:** Automated incident narratives can support security reporting, auditing, and post-incident review.
14. **Policy-Aware Decision Support:** Security recommendations can be constrained by organizational policies and governance requirements.
15. **Adaptive Cyber Defense:** Continuous learning from new incidents and updated knowledge sources can improve threat-analysis capabilities over time.

Disadvantages

1. **Hallucination Risk:** LLMs may generate plausible but unsupported security conclusions.
2. **Prompt-Injection Threats:** Malicious content embedded in logs or retrieved documents may attempt to manipulate model behavior.
3. **High Computational Requirements:** Large models can require substantial processing resources, particularly for private deployment.
4. **Sensitive Data Exposure:** Poorly designed implementations could expose confidential security telemetry to unauthorized systems.

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0706010

5. **Model Dependency:** Organizations may become dependent on specific models, vendors, or specialized AI infrastructure.
6. **Limited Cybersecurity Understanding:** General-purpose LLMs may misunderstand highly specialized infrastructure or attack scenarios.
7. **Retrieval Errors:** RAG systems may retrieve irrelevant, incomplete, outdated, or manipulated information.
8. **High Deployment Complexity:** Integrating LLMs with cloud telemetry, SIEM systems, threat intelligence, IAM, and governance systems can be technically challenging.
9. **Operational Cost:** Hosting, maintaining, fine-tuning, monitoring, and updating LLM infrastructure can be expensive.
10. **False Conclusions:** Incorrect AI interpretation can cause analysts to investigate the wrong systems or prioritize the wrong threats.
11. **Latency:** Large-model inference and multi-stage retrieval can increase response time for real-time security analysis.
12. **Model Drift:** Changes in threats, cloud architectures, terminology, and organizational policies can reduce analytical accuracy.
13. **Governance Complexity:** Sovereign deployments require strict controls over data residency, model access, auditing, and administrative privileges.
14. **Adversarial Manipulation:** Attackers may intentionally create security events designed to confuse or mislead AI-based analysis.
15. **Human Oversight Still Required:** High-impact security decisions cannot safely depend entirely on autonomous LLM reasoning and require appropriate analyst validation.

IV. RESULTS AND DISCUSSION

The proposed **Large Language Model-Based Intelligent Cyber Threat Analysis for Sovereign Cloud Infrastructure** was evaluated as an intelligent security architecture designed to improve cyber-threat identification, contextual analysis, incident interpretation, and response support within sovereign cloud environments. The framework combines large language models (LLMs) with conventional security monitoring, threat intelligence, log analytics, identity information, cloud telemetry, and policy repositories. The evaluation indicates that LLM-based threat analysis can significantly improve the ability of security teams to interpret large volumes of heterogeneous cybersecurity information while maintaining the governance requirements associated with sovereign infrastructure. Traditional security systems generally generate alerts independently, requiring analysts to manually correlate authentication events, network anomalies, vulnerability information, configuration changes, and threat-intelligence indicators. The proposed approach instead uses an LLM-driven analytical layer to correlate these heterogeneous signals and construct a contextual representation of potential incidents. For example, an isolated failed authentication event may have limited significance, whereas a sequence involving repeated authentication failures, successful access from an unusual location, privilege escalation, access to sensitive resources, and abnormal data transfers can represent a potentially coordinated intrusion. The LLM-based analytical layer can summarize such sequences, identify relationships among events, classify potential attack techniques, and produce risk-oriented explanations for analysts. The results suggest that this contextual reasoning capability is particularly valuable for sovereign clouds because such environments often contain sensitive governmental, public-sector, defense, financial, healthcare, or nationally important information requiring strict security controls. Rather than sending sensitive logs to external AI services, the framework assumes deployment of the language model within a controlled sovereign environment or through an approved isolated inference infrastructure. This architecture preserves greater control over data residency, model access, security policies, and auditability. The results therefore indicate that LLMs can complement conventional detection technologies by transforming fragmented security events into understandable and actionable threat intelligence while maintaining stronger control over sensitive sovereign data.

The second major finding concerns **improved threat correlation, incident interpretation, and analyst efficiency**. Sovereign cloud infrastructures generate enormous quantities of operational and security telemetry from virtual machines, containers, Kubernetes clusters, identity providers, APIs, firewalls, endpoint systems, databases, storage services, and network infrastructure. Conventional rule-based systems are effective at identifying known patterns but may produce large numbers of alerts that require manual prioritization. The proposed LLM-based framework addresses this challenge through semantic analysis and contextual correlation. Security events are normalized and enriched with asset criticality, identity context, vulnerability information, historical behavior, and threat-intelligence information before being presented to the analytical model. The LLM can then generate incident summaries, identify potentially related events, explain attack progression, and recommend investigation priorities. This capability reduces the cognitive burden on security analysts, particularly when incidents involve multiple infrastructure layers. The results further demonstrate the value of retrieval-augmented generation, where the LLM retrieves information from trusted sovereign sources such as internal security policies, vulnerability databases, incident histories, architecture documentation, and

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0706010

approved threat-intelligence repositories before generating an analytical response. This reduces dependence on the model's internal knowledge and allows its output to reflect organization-specific security requirements. Importantly, retrieval mechanisms also provide a pathway for evidence-based analysis because generated conclusions can be connected to underlying events and authoritative information. The framework can therefore distinguish between an observed fact, an inferred threat, and a recommended response. This distinction is essential for security operations because an LLM should not present an uncertain inference as a confirmed incident. The results also suggest that analyst productivity can be improved when LLMs perform repetitive activities such as log summarization, event grouping, incident-description generation, attack-path explanation, and initial threat classification. Human analysts can consequently focus on higher-value activities such as validating critical findings, conducting forensic investigations, improving policies, and coordinating responses. However, the evaluation also highlights the importance of confidence estimation and human verification because language models may produce plausible but incorrect explanations when evidence is incomplete.

A further result concerns **sovereignty, privacy, governance, and security control**. Sovereign cloud infrastructure differs from conventional public-cloud environments because organizations may require strict control over where data is stored, processed, and accessed and may impose additional requirements concerning national regulations, organizational policies, and trusted infrastructure. The proposed framework addresses these requirements by keeping sensitive security telemetry inside the sovereign security boundary. Instead of transferring raw logs or confidential infrastructure information to externally hosted LLM services, the architecture can employ locally hosted or sovereign-approved models. Access to the model is controlled through identity and authorization mechanisms, while security prompts, retrieved documents, model outputs, and response actions can be logged for auditing. The results indicate that this architecture substantially improves governance because administrators can determine which information the model can access and which actions it is authorized to recommend or execute. Data classification mechanisms can additionally prevent highly sensitive information from being unnecessarily exposed to the language model. Encryption, secure key management, access controls, network segmentation, and confidential-computing mechanisms can further strengthen the AI processing environment. Another significant finding is that LLMs should not independently make unrestricted security decisions. Instead, the framework uses a policy-governed architecture in which model recommendations are evaluated against predefined security rules and organizational constraints. For instance, an LLM may recommend isolating a suspected compromised workload, but an enforcement engine determines whether that action is permitted based on workload criticality and operational policy. This separation between reasoning and enforcement improves safety and accountability. The results also reveal limitations. Sovereign deployment can require significant computational resources, especially for large models. Model updates, hardware requirements, inference latency, and specialized security expertise can increase operational complexity. Furthermore, language models remain vulnerable to hallucination, prompt injection, data poisoning, and manipulation through malicious log content. Therefore, LLM outputs must be treated as security intelligence rather than unquestionable truth. The strongest architecture is one in which model-generated conclusions are supported by retrieved evidence, deterministic security controls, confidence thresholds, and human approval for high-impact decisions.

Overall, the results demonstrate that **LLM-based intelligent cyber-threat analysis can strengthen sovereign cloud security when implemented within a controlled, evidence-driven, and policy-governed architecture**. The framework provides a bridge between conventional security analytics and advanced semantic reasoning. Conventional systems remain responsible for collecting telemetry, enforcing policies, identifying deterministic indicators, and executing security controls, while the LLM contributes contextual interpretation, event correlation, natural-language explanation, threat summarization, and analyst assistance. This complementary design is more practical than attempting to replace existing security infrastructure with an LLM. The framework is particularly useful for complex incidents where security evidence is distributed across different systems and requires human-like contextual interpretation. It can help analysts understand how an incident developed, which assets may be affected, what attack techniques are likely involved, and which response actions should be considered. At the same time, the results indicate that sovereign deployment introduces a unique balance between intelligence and control. Organizations must obtain the benefits of advanced AI while preserving data sovereignty, regulatory compliance, infrastructure independence, auditability, and operational resilience. Model performance should therefore be evaluated not only through conventional language-model metrics but also through cybersecurity measures such as detection accuracy, false-positive rate, incident-triage time, evidence-groundedness, response latency, analyst workload reduction, and policy-compliance rate. Security-specific evaluation should also test the model against adversarial prompts and malicious telemetry. Despite these challenges, the overall findings support the feasibility of LLM-based cyber-threat analysis as an important component of sovereign cloud security. By combining locally controlled AI, trusted knowledge retrieval, security telemetry, policy enforcement, and human oversight, the framework can improve threat visibility and accelerate security operations without compromising the fundamental sovereignty requirements of sensitive cloud environments.

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0706010

V. CONCLUSION

The development of a **Large Language Model-Based Intelligent Cyber Threat Analysis for Sovereign Cloud Infrastructure** represents an important step toward integrating advanced artificial intelligence with security architectures designed for environments where data sovereignty, regulatory control, operational independence, and national security are critical. Sovereign clouds increasingly support sensitive workloads that cannot always rely on conventional external cloud services or unrestricted data-processing models. Consequently, security systems operating in these environments must simultaneously provide advanced threat-detection capabilities and strict control over sensitive information. The proposed framework addresses this requirement by placing an LLM-based analytical layer within a controlled sovereign security architecture. Instead of replacing established security technologies, the framework enhances them through semantic reasoning and contextual analysis. Security events generated by identity systems, network infrastructure, workloads, applications, APIs, databases, and cloud platforms can be consolidated, enriched, and analyzed by the language model. The model can identify relationships among events, summarize complex incidents, classify potential threats, explain suspicious behavior, and assist analysts in determining appropriate investigation paths. This capability is particularly important because modern attacks rarely occur as isolated events. Attackers frequently move across identities, applications, networks, and infrastructure components, leaving fragmented indicators that may not be recognized by individual security tools. The proposed LLM-based approach helps connect these fragments into coherent incident narratives. The conclusion from the evaluation is therefore that language models can significantly enhance the analytical layer of sovereign cybersecurity when their capabilities are combined with reliable telemetry, trusted knowledge sources, and enforceable security policies.

A central conclusion is that **sovereignty must remain a foundational design principle rather than an afterthought**. The use of external generative AI services can create risks involving sensitive log exposure, data residency, intellectual property, regulatory compliance, and unauthorized access. The proposed architecture therefore emphasizes locally controlled or sovereign-approved LLM deployment. By keeping sensitive security information within the organization's trusted environment, the framework provides greater control over how cybersecurity data is processed and retained. This approach also enables organizations to establish customized access policies, model governance procedures, audit mechanisms, and security boundaries. Data minimization can ensure that only relevant information is presented to the model, while anonymization or pseudonymization can further reduce unnecessary exposure of sensitive identifiers. Encryption and secure infrastructure provide additional protection for both model inputs and outputs. Another important conclusion is that sovereignty extends beyond physical data location. Organizations must also consider who controls the model, where model updates originate, who can modify prompts and retrieval sources, and how security decisions are audited. Dependence on external model providers may create strategic or operational risks even when data residency requirements are technically satisfied. A sovereign LLM architecture can consequently provide greater institutional control over the complete AI-security lifecycle. However, this independence introduces additional responsibilities. Organizations must manage model infrastructure, security updates, performance monitoring, training or fine-tuning, hardware capacity, and vulnerability management. Therefore, sovereign AI should be viewed as a broader governance and infrastructure strategy rather than merely a decision to host an LLM locally.

The findings further establish that **LLM reasoning should be constrained by evidence and deterministic controls**. Although language models provide powerful natural-language reasoning capabilities, they can generate inaccurate or unsupported statements. In cybersecurity, such errors can result in inappropriate incident escalation, unnecessary workload isolation, missed threats, or incorrect remediation actions. The proposed framework addresses this limitation through retrieval-augmented analysis, confidence assessment, evidence association, and policy-based response mechanisms. An LLM can interpret security evidence and produce a recommendation, but critical decisions should be validated against trusted data and organizational security policies. This creates a layered architecture in which the LLM functions as an intelligent analyst rather than an unrestricted autonomous administrator. Such separation is particularly important in sovereign environments where infrastructure availability and security controls may have national or organizational significance. Explainability is another essential component. Security teams must be able to determine why the model classified an event as suspicious and which evidence influenced its recommendation. An effective LLM-based system should therefore identify supporting telemetry, relevant threat-intelligence information, affected assets, confidence levels, and potential alternative explanations. This improves transparency and enables human analysts to challenge or validate model conclusions. The framework consequently supports a human-AI collaboration model in which AI performs high-volume analytical work while humans retain authority over consequential decisions. This arrangement can improve both efficiency and accountability. The conclusion is that responsible LLM integration requires carefully defined boundaries between analysis, recommendation, and enforcement.

Finally, the proposed framework demonstrates the potential of LLMs to transform **sovereign cloud security operations from fragmented alert management into contextual and intelligent cyber defense**. By integrating

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2024.0706010

language models with cloud telemetry, threat intelligence, security policies, identity information, and incident-response mechanisms, organizations can obtain a more comprehensive understanding of complex cyber events. The system can reduce the time required to interpret large quantities of security information, assist with incident prioritization, generate understandable reports, and improve communication between technical and managerial stakeholders. Nevertheless, the framework should not be considered a complete replacement for established cybersecurity mechanisms. Firewalls, endpoint protection, vulnerability management, identity security, encryption, network segmentation, secure configuration, intrusion detection, and access control remain essential. The LLM provides an additional intelligence layer that helps security teams understand and coordinate these capabilities. The long-term value of the framework will depend on its ability to maintain accuracy, security, transparency, and sovereignty while adapting to emerging threats. Future implementations should therefore emphasize rigorous benchmarking, adversarial testing, secure model lifecycle management, evidence-grounded generation, and continuous policy validation. Overall, the study concludes that LLM-based cyber-threat analysis can provide substantial benefits to sovereign cloud infrastructures when deployed within a secure and governed architecture. The combination of advanced language reasoning, trusted security data, sovereign deployment, policy enforcement, and human oversight provides a practical foundation for next-generation cyber defense. Such an architecture can help sovereign-cloud operators improve threat intelligence, accelerate incident investigation, strengthen security governance, and maintain greater control over sensitive digital infrastructure.

VI. FUTURE WORK

Future research should focus on developing **domain-specific and sovereign-aware language models** that are optimized for cybersecurity analysis rather than relying exclusively on general-purpose LLMs. General models possess broad language capabilities, but they may not fully understand the specialized terminology, architectures, regulations, operational practices, and threat landscapes associated with sovereign cloud environments. Future models could be trained or fine-tuned using carefully governed cybersecurity datasets containing security logs, cloud events, vulnerability information, attack patterns, incident reports, policy documents, and threat-intelligence material. However, training data must be carefully curated to avoid incorporating sensitive information or biased operational assumptions. Privacy-preserving learning techniques could allow organizations to improve models without directly sharing confidential telemetry. Federated learning, secure aggregation, and privacy-enhancing computation could be investigated for collaborative model development across sovereign organizations. Future research should also investigate smaller domain-specific models that can provide high-quality security reasoning with lower computational requirements. Such models may be more practical for sovereign environments with limited infrastructure resources or strict requirements for isolated processing. Model compression, quantization, knowledge distillation, and specialized inference optimization could reduce hardware requirements while maintaining acceptable cybersecurity performance. Another important research direction is multilingual security intelligence. Sovereign infrastructures may operate across multiple languages, and future models should be capable of analyzing security alerts, incident reports, policies, and threat intelligence across languages without losing technical meaning. These developments could create more accessible and locally controlled AI-security capabilities while reducing dependence on externally hosted general-purpose models.

A second major research direction involves **secure retrieval-augmented generation and evidence-grounded cyber reasoning**. Future systems should move beyond simple retrieval of documents and develop structured security knowledge graphs connecting identities, assets, vulnerabilities, network relationships, attack techniques, events, and historical incidents. Such graphs could provide the LLM with a continuously updated representation of the sovereign cloud environment. Combining knowledge graphs with retrieval-augmented generation could improve the model's ability to identify attack paths and explain relationships among seemingly unrelated events. Research should establish mechanisms for verifying every security conclusion against available evidence and identifying situations where sufficient evidence does not exist. Confidence calibration should be investigated so that the model can distinguish high-confidence findings from speculative hypotheses. Future architectures could require the LLM to provide evidence references, affected assets, reasoning categories, confidence scores, and alternative explanations for each security recommendation. This would reduce hallucination and improve auditability. Research should also investigate resistance to prompt injection originating from malicious logs, compromised applications, or attacker-controlled data. Since security telemetry can contain arbitrary text, an attacker could attempt to manipulate the LLM by embedding instructions within log messages or other retrieved content. Future systems therefore need strict separation between trusted system instructions and untrusted telemetry. Content sanitization, retrieval-source trust scoring, instruction isolation, and output verification should be explored. These capabilities would be essential for building LLM-based security systems that can safely operate in adversarial environments.

Future work should also investigate **autonomous and semi-autonomous cyber response for sovereign cloud infrastructure**. Current approaches can use LLMs primarily for analysis and recommendation, but future systems

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0706010

could support carefully bounded automated response. An intelligent security agent might investigate suspicious events, retrieve additional evidence, execute predefined diagnostic procedures, assess attack severity, and recommend containment actions. However, autonomous response should operate within strict policy boundaries. Future research could develop graduated autonomy levels in which low-risk activities are fully automated, moderate-risk actions require validation, and high-impact actions require explicit human approval. Reinforcement learning could potentially optimize response strategies by considering security improvement, service availability, operational cost, and recovery time. Digital twins and cyber ranges could provide safe environments for training and evaluating such autonomous agents without affecting production infrastructure. Research should measure whether AI-driven response reduces mean time to detect, mean time to investigate, and mean time to respond while maintaining service availability. Additional work should explore coordination between LLM agents and existing security orchestration, automation, and response platforms. Rather than allowing an LLM to directly control infrastructure, a secure orchestration layer could validate AI recommendations against predefined playbooks before execution. This architecture would combine the flexibility of AI with the predictability of deterministic security automation. Future research should also examine failure modes, including incorrect isolation, excessive remediation, model unavailability, compromised agents, and conflicting security policies.

Finally, future research should establish **standardized evaluation, governance, and interoperability frameworks** for LLM-based sovereign cyber defense. Current LLM evaluation often emphasizes language-generation quality, but cybersecurity systems require substantially different measures. Future benchmarks should evaluate detection accuracy, threat-classification performance, evidence-groundedness, hallucination rate, false-positive rate, false-negative rate, incident-triage time, response latency, analyst workload reduction, computational cost, policy compliance, and resilience against adversarial manipulation. Standardized datasets should include realistic sovereign-cloud attack scenarios and diverse infrastructure configurations. Researchers should also develop governance frameworks addressing model ownership, data residency, auditability, update management, accountability, access control, and secure model retirement. Interoperability is another important requirement. Future LLM security platforms should integrate with Kubernetes security tools, cloud security platforms, SIEM systems, identity providers, vulnerability scanners, threat-intelligence platforms, security orchestration systems, and existing governance frameworks. Open interfaces and standardized data representations could reduce vendor dependence and support deployment across heterogeneous sovereign environments. Long-term research should additionally evaluate the organizational implications of AI-assisted security, including analyst training, changes in security operations-center workflows, human-AI trust, and accountability for AI-supported decisions. Ultimately, the goal should be to establish a **trustworthy sovereign AI cyber-defense ecosystem** in which language models provide powerful contextual intelligence without weakening security, privacy, sovereignty, or human accountability. Through secure model engineering, evidence-grounded reasoning, controlled autonomy, standardized benchmarking, and strong governance, future LLM-based systems can evolve from analytical assistants into reliable components of resilient sovereign cloud cybersecurity.

REFERENCES

1. Belal, M. M., & Sundaram, D. M. (2022). Comprehensive review on intelligent security defences in cloud: Taxonomy, security issues, ML/DL techniques, challenges and future trends. *Journal of King Saud University - Computer and Information Sciences, 34*(10), 9102–9131. https://doi.org/10.1016/j.jksuci.2022.08.035
2. Gummadi, V. P. K. (2021). Secure API lifecycle management: Integrating MuleSoft Secrets Manager for enterprise data protection. International Journal of Intelligent Systems and Applications in Engineering, 9(4), 537-540.
3. Rella, B. P. R., Chakraborty, A., Shah, S. B., Ghosh, H., Gautam, S., Dhirwani, B. A., ... & Mutyam, N. (2024). AI-engine-based acceleration for high-performance programmable system-on-chip designs. Journal of Computational Analysis and Applications, 32(1).
4. Jayaraman, S., Rajendran, S., & P, S. P. (2019). Fuzzy c-means clustering and elliptic curve cryptography using privacy preserving in cloud. International Journal of Business Intelligence and Data Mining, 15(3), 273-287.
5. Beeram, S. (2024). AI-driven intelligent traffic management in Microsoft Azure: Predictive routing for resilient cloud applications. International Journal of Advanced Engineering Science and Information Technology, 7(4), 14534–14538.
6. Challa, R. (2023). Reliability Engineering for Zero-Downtime Integration of HPC Systems into Regulated Environments. International Journal of Research and Applied Innovations, 6(6), 10082-10092.
7. Vani, M., & Dadlani, D. (2023). Smart home – "Aashraya" IoT automation system. International Journal of Computer Technology and Electronics Communication, 6(1), 6393–6403.
8. Alam, A., Gazi, M. S., Abdullah, S. M., Tasnim, M., Himeluzzaman, M., Nabil, M. A., ... & Akter, S. (2021). Quantum-resilient federated intrusion detection: A hybrid quantum-classical framework for safeguarding US

|ISSN 2349-3216| Volume 7, Issue 6, November-December 2024| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2024.0706010

critical infrastructure in the post-quantum era. International Journal of Advances in Signal and Image Sciences, 7(1), 57–72.

9. Chaba, A. (2021). API-driven enterprise commerce architecture for composable digital ecosystems. International Journal of Engineering & Extended Technologies Research, 3(6), 4082–4086.
10. Mathew, A., & Romasco, L. (2024). Forensic Investigation of Artificial Intelligence Systems. Research Updates in Mathematics and Computer Science Vol. 4, 154-164.
11. Anand, L. (2022). Integrating Kubernetes Microservices with Privileged Access Security and Real-Time Fraud Detection for Modern Enterprise Systems. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 5(5), 7453-7461.
12. Koganti, H. (2024). The computational complexity of hybrid algorithms: When branch-and-bound meets machine learning heuristics. International Journal of Research and Applied Innovations (IJRAI), 7(4), 11184–11190.
13. Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. Indian Journal of Science and Technology, 9, 44.
14. Padmanabham, S. (2023). Secure API gateway architecture for open banking. International Journal of Computer Technology and Electronics Communication, 6(6), 8166–8174.
15. Sabin Begum, R., & Sugumar, R. (2019). Novel entropy-based approach for cost-effective privacy preservation of intermediate datasets in cloud. Cluster Computing, 22(Suppl 4), 9581-9588.
16. Vemireddy, S. (2023). Building resilient enterprise platforms using event-driven and distributed computing models. International Journal of Research and Applied Innovations (IJRAI), 6(6), 10106–10112.
17. Hoque, M. J., Hasan, M. M., Khatun, M. M., Akter, F., & Mohammad, A. R. (2021). Impact of COVID-19 on Consumer Buying Behavior During COVID-19 Pandemic Using Data Analytics. Journal of Business and Management Studies, 3(2), 296-307.
18. Vimal Raja, G. (2024). Intelligent data transition in automotive manufacturing systems using machine learning. International Journal of Multidisciplinary and Scientific Emerging Research, 12(2), 515-518.
19. Yepuri, V. K., Polamarasetty, V. K., Donthi, S., & Gondi, A. K. R. (2023). Containerization of a polyglot microservice application using Docker and Kubernetes.arXiv preprint arXiv:2305.00600
20. Bellundagi, M. (2023). Design of an Intelligent Clinical Decision Support System Using Machine Learning Techniques. International Journal of Research and Applied Innovations, 6(6), 10075-10081.
21. Raja, G. V. (2020). Metadata gets a makeover: The machine learning approach. International Journal of Computer Technology and Electronics Communication, 3(6), 2900-2903.
22. Abd-Rouf, M. S. K., Adigun, P. O., Alalade, E. O., Oyekanmi, T. T., Faniyi, A. J., Oladapo, B., Awopejo, T. E., Adegoke, O. S., Jamiu, A., Michael, O. B., Obisesan, A., Ajala, S., Adekanye, M. A., Yambali, P. M., & Abd-Rouf, A. B. (2024). From molecular profiling to predictive algorithms: A conceptual machine-learning framework for mechanism-informed therapy selection in multidrug-resistant cancer. International Journal of Science, Research and Technology (IJSRAT), 7(3), 12085–12101.
23. Gopinathan, V. R. (2023). Modernizing Enterprise Applications Using Intelligent API Frameworks Machine Learning and Cloud Native Computing. International Journal of Science, Research and Technology, 6(5), 10690-10697.
24. Bandaru, P. K. (2021). Leveraging hardware-in-the-loop simulation for continuous automotive software testing. International Journal of Engineering & Extended Technologies Research (IJEETR), 3(5), 3730–3735.
25. Soundappan, S. J. (2023). Designing Intelligent Enterprise Platforms Using Machine Learning Driven API Engineering and Cloud Native Security. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(4), 9074-9081.
26. Rahul Reddy Bandhela, RamMohan Reddy Kundavaram. (2023). A Comparative Study on Neural Network Architectures for Image Recognition Applications . Journal of Computational Analysis and Applications (JoCAAA), 31(1), 1334–1342. Retrieved from https://www.eudoxuspress.com/index.php/pub/article/view/3566
27. Sivakumer, D. (2024). From posts to profits: A systematic review on how social media influencers shape brand communication and success. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(1), 10042–10055.
28. Raja, G. V. (2023). AI-Driven Cloud-Native Enterprise Systems Leveraging Kubernetes, DevSecOps, and Predictive Analytics. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(2), 7925-7929.
29. Rajula, A. (2022). Cloud-based virtual patient engagement with intelligent scheduling and secure document management. *International Journal of Future Innovative Science and Technology (IJFIST)*, 5(5), 9245.
30. Chy, M. S. K., Abdullah, S. M., Nabil, M. A., Alam, A., Himeluzzaman, M., Onik, T. A., ... & Khan, H. A. (2022). QShield-NS: A Variational Quantum Machine Learning Model for Zero-Day Cyber Threat Detection in National Security Systems. International Journal of Future Innovative Science and Technology (IJFIST), 5(5), 9283.
31. Wazid, M., Das, A. K., Chamola, V., & Park, Y. (2022). Uniting cyber security and machine learning: Advantages, challenges and future research. *ICT Express, 8*(3), 313–321. https://doi.org/10.1016/j.icte.2022.04.007