

LLM-Driven Autonomous Auditors on Blockchain: A New Framework for Verifiable Cloud Security

Dr. Nidal Al Said

Associate Professor, College of Mass Communication, Ajman University, Ajman, UAE

Abstract: The high rate of cloud-native infrastructures and distributed applications development has posed challenges to the assurance of continuous, reliable and verifiable security never before witnessed. Conventional methods of auditing are intensive in manual inspection, centralised monitoring tools and periodical compliance testing and in most cases they do not meet the dynamic nature of cloud environments. The recent developments in Large Language Models (LLMs) have shown good reasoning, anomaly detection, and policy-interpretation, which have contributed to the creation of a new category of autonomous agents. Similarly, blockchain technology provides distributed, cryptographically verifiable, immutable ledgers which are suitable to have transparent audit trails. This paper proposes a single architecture of autonomous auditor with the support of LLM deployed on the blockchain networks to guarantee real-time security, policy adherence and tamper-proof verifiability in the cloud ecosystems. The suggested architecture combines smart contracts in order to enforce the rules, decentralized logs in order to provenance, LLM-based agent in order to interpret and make decisions, and a reinforcement learning loop in order to self-optimize the auditor. The outcomes of experimental reasoning are marked by a higher level of consistency in the audits, a decrease in human overhead, and a high degree of resistance to manipulation. The framework opens a way to sovereign AI auditors, who can provide unbiased, constant, and verifiable cloud security.

Keywords: LLM Auditors; Blockchain Security; Autonomous Agents; Cloud Compliance; Verifiable Computing; Smart Contracts; Zero-Trust Cloud; Decentralized Audit Trails; AI Governance; Reinforcement Learning; Cloud Forensics; Secure Multi-Agent Systems.

I. INTRODUCTION:

The current cloud-based infrastructures are marked by the existence of multi-tenant designs, containerization organizations, virtualized networks, and distributed compute nodes, which have resulted in increased complexity in their operations than before. Conventional security auditing systems, including regular compliance assessments, one-off rule-based scanners, and control reviews conducted by human beings, find it difficult to offer continuous coverage in milliseconds-scale changing environments. Ensuring manual verification is made more difficult by the emergence of automated DevSecOps pipelines and the evanescent cloud resources. Subsequently, organizations tend to demand real-time, transparent and reliable security controls that can identify misconfigurations, policy abuse and suspicious behavior at scale.

At the same time, the new trends in the field of Large Language Models (LLMs) have altered the possibilities of autonomous systems in the security domain. The LLMs are good at understanding complicated policies, log analysis, summarization of the telemetry, detecting anomalies, explaining attacks, and generating remedies. With the addition of agent architectures, they become autonomous auditors that can assess system states and activate investigations as well as prescribe corrective controls without human intervention.

Another important component that has been brought about by blockchain technology is audit trails that are verifiable and

cannot be manipulated. Unchangeable registries give a cryptographically secure record of all audit processes, model choices, data flows and compliance checks. Thus, the collaboration of the auditors with LLM and the verification with blockchain is a unique opportunity to redefine cloud security auditing.

The current paper suggests a new universal and broad-scale framework in which the role of the LLM agents is that of the autonomous auditor, and blockchain smart contracts are used as execution enforcement agents and integrity assurances. The architecture that results is designed so that all the audit interactions, decision, and remedial activities are recorded unalterably, facilitating high levels of transparency and credibility by the regulators, administrators, and cloud service consumers.

II. BACKGROUND AND MOTIVATION

Cloud workloads are dynamic and in the context of modern environment, resources come and go within a short time and configurations are constantly changing due to continuous pipelines of automation. Conventional security auditing systems do not work well in such an environment due to several factors such as attempts to identify misconfigurations are normally identified too late, the tools used are usually fragmented and human analysts are very much relied upon to interpret as well as make decisions. The centralized audit logs also do not ease the issue of trust as they can be easily manipulated, insider threats,

and single points of failure. These shortcomings inform the fact that without moving to the stronger approach to audit, there was a need to create more resilient, real-time, and trustful forms of audit.

At the same time, the creation of Large Language Models (LLMs) has introduced a new generation of autonomous actors possessing good reasoning, in-depth contextual understanding, and intelligent policy-making skills. Based on the agent-based design, LLMs have the ability to scan the cloud infrastructure on their own, comprehend multi-faceted compliance standards, such as NIST and ISO 27001, analyze access control policies, and summarize risks as they evolve with high accuracy. They are cognitive in nature and are able to dynamically respond to new actions by the threats without necessarily employing hard-determined rule-based systems and this qualifies them very well to be employed in the complex and highly changing cloud environments.

The other dimension that puts a critical perspective on trustful auditing is a blockchain technology that comprises of decentralized and storage-resistant audit logs storage. There is cryptographic hashing to make sure that the recorded events are immutable to leave clear and verifiable evidence trails, which can be comprehensively viewed individually. Smart contracts facilitate such properties by enabling the possibility of enforcing audit rules, where the enforcement is automated, audit data access control and workflow validation is deterministic. The combination of the autonomous agency of LLM and the verification of blockchains directly overcomes the limitations of the old, centralized system of auditors and creates a more reliable, more transparent, and safer auditor ecosystem.

III. PROPOSED FRAMEWORK

LLM-Driven Autonomous Auditors on Blockchain

The proposed system integrates the audit agent of LLM, blockchain-based architecture, decentralized data provenance, and real-time streamlines cloud telemetry into a unified architectural model, which can be used to prove and tamper-proof audit cloud. This alone system is capable of 24 hour monitoring, intelligent analysis and evidence creation that can not be altered, and this breaks the restrictions of the conventional cloud audit systems.

Attack Events on Kippo Honeypot




Figure 1. Distribution of attack events observed on the Kippo honeypot during monitoring.

At the architectural level the system is designed based on four inter-linked components that work in a synergistic manner. The autonomous auditors powered by LLM can be viewed as the brain of this system, as they are capable of performing a variety of complex reasoning, detection of anomalies, policy interpretation, log analysis, and security decision-making. A blockchain layer supporting these agents is in place to store the audit evidence that aids compliance enforcement through the execution of smart contracts and offers trusted, decentralized, foundational support to the work of auditors. In addition to these elements, there is a cloud telemetry interface, which constantly gathers logs, configuration variants, resource conditions, access patterns, and system occasions of various cloud settings. To enhance sustained performance, a reinforcement learning optimization loop allows auditors to improve their behavior through time based on the feedbacks about the accuracy of their detections, false positives, effectiveness of incident response, as well as the quality of policy interpretation.

Table 3.1. Core components of the proposed LLM-Blockchain autonomous auditing framework and their primary responsibilities.

Component	Description	Key Responsibilities
LLM-Driven Autonomous Auditor	AI agent capable of reasoning, interpreting policies, detecting anomalies.	Real-time risk assessment, configuration analysis, policy compliance evaluation.
Cloud Telemetry Interface	Collects logs, resource states, IAM activity, configuration drifts.	Stream ingestion, normalization, event classification.
Blockchain Layer	Immutable, decentralized evidence ledger.	Audit event verification, tamper-proof storage, transparency.
Smart Contracts	Deterministic policy enforcement and validation mechanisms.	Access control rules, compliance logic, automated remediation triggers.
Reinforcement Learning Module	Performance feedback loop for the LLM auditor.	Optimizes detection accuracy, reduces false positives, adapts to new threats.

The autonomous auditing begin with the LLM agent that is again and again processing telemetry data to get an idea of the security posture of the system and be able to detect any change in the configuration of the system or a change in the behavioral pattern. Upon detecting an anomaly, misconfiguration or policy violation by the agent, it generates a verifiable audit event and sends it to

the blockchain. Smart contracts in turn verify the authenticity of the occurrence and eliminate the registering of fraudulent evidence. Once validated, the information contained in the audit is permanently written and the smart contract may be able to take corrective actions automatically or warn the interested parties, which could be regarded as a useful method of responding effectively and safely to the threat as early as possible.

In this architecture, the use of smart contracts is significant when it comes to compliance requirement enforcing. These policies are identity and access management controls, cryptography requirements, encryption requirements, network segmentation policies and governance requirements which are measurable and encoded. The LLM agent is applied to these contracts to make sure that they satisfy policy requirements and produce evidence-based outputs, which are transparently stored on-chain and provide consistency and determinism in the enforcement of compliance checks.

In addition, the framework provides the decentralized logs storage and hash-anchored evidence trails to facilitate the fact that all audit artifacts are preserved. The snapshots of the system, and summaries of system configuration and reasoning, and diagnostic reports are hashed with cryptography and stored in the blockchain that does not allow them to be altered. This will ensure non-repudiation besides ensuring that regulators, auditors and internal security teams have a reliable and verifiable audit trail that they can consult at any specific time.

IV. REINFORCEMENT LEARNING FOR AUTITOR SELF OPTIMIZATION

The LLM auditor leverages reinforcement learning (RL) to continuously refine performance. Feedback signals include detection accuracy, precision of alerts, time-to-diagnosis, correctness of policy interpretations, and successful smart-contract interactions. Over time, the agent optimizes its reasoning strategies, anomaly scoring thresholds, and remediation suggestions, creating an increasingly effective auditor.

V. SECURITY MODEL AND THREAT ANALYSIS

The suggested framework focuses on a number of the most serious threats encountered in the contemporary clouds through the integration of intelligence based on LLM with the verification supported by blockchain. It efficiently controls insider threats, log manipulation attacks, configuration drift, privilege escalation and non-compliance. The system attempts to ensure the absence of historical tampering with audit events due to immutable storage in the blockchain and the continuous monitoring and swift identification of abnormal behaviors, offered by the agents of the LLM.

In spite of these advantages, there is also a new risk as a result of the introduction of auditors who are driven by LLM. Unless constrained, LLAs may hallucinate, false interpretations of ambiguous telemetry, or false recommendations. In response to these vulnerabilities, the framework includes multi-agent

validation, cryptographic validation of audit events, consensus-based decision controls and smart contract mechanisms which restrict an agent to undertake high impact actions unless they are deterministically checked. The following safeguards are needed to make sure that the probabilistic character of the LLM reasoning would not affect the overall reliability of the system.

Table 5.1. Major cloud-security threats and the mechanisms within the proposed framework that address each risk.

Threat Category	Example Attack	Framework Defence Mechanism
Insider Threats	Log manipulation, access abuse	Blockchain immutability; smart-contract validated actions; LLM anomaly detection
Configuration Drift	Unapproved security-group changes	Continuous telemetry monitoring + LLM policy comparison
Privilege Escalation	IAM role misuse	Smart contracts enforcing strict role logic; cross-agent validation
Telemetry Poisoning	Injecting deceptive logs	Input sanitization, multi-model verification, cryptographic signatures
Consensus/Blockchain Attacks	51% attack attempts	Permissioned chain governance; redundant validator nodes
LLM Hallucination Risk	False interpretation of evidence	Smart-contract deterministic validation; multi-agent cross-checking

Furthermore, the architecture considers the adversarial attacks on either the blockchain infrastructure or the models. The threat actors can also seek to do prompt injection, telemetry poisoning, or consensus attack on the chain. The system helps to address these risks with redundancy in the validator nodes, cross-checking of the results with the help of numerous independent models, the strict sanitization and filtering of telemetry inputs, and effective policy of blockchain governance that helps to prevent the attacks on the consensus-layer. Combined, these levels of defense make the whole system of auditing more resilient to advanced attackers.



Figure 1: The NIST Cybersecurity Framework

VI. IMPLEMENTATION OF CONSIDERATIONS

The choice of blockchain is one of the critical factors that influence the performance and reliability of the suggested auditing framework. Although blockchains used publicly are highly decentralized, they present trigger issues like lack of throughput, increased transaction costs, and loss of control over governance. Conversely, the necessary scalability, selective access and customizable management required in enterprise cloud auditing are found by private or permissioned blockchains. The properties allow permissioned chains to be more appropriate in a secure management of high-frequency audit events, compliance logic enforcement, and for providing organizational oversight.

Privacy and data protection parameters also have to be taken strictly when integrating LLMs in the auditing process. Any interactions of the model should adhere to encryption frameworks, organizational data-handling policies, and confidential computer usage procedures that reduce the exposure of confidential information. To add additional benefits, a sensitive telemetry or configuration data can be pre-processed with a carefully managed differential privacy, zero-knowledge proofs, or other privacy preserving computation methods and then analyzed by the LLM or stored in the blockchain. This makes sure that the system is effective and in compliance with expectations on regulations.

The framework will be placed to back cross-cloud functionality by means of joining with Kubernetes clusters, virtual machine environments, serverless functions, and an assortment of cloud-native telemetry API. Its cloud agnostic connectors can be deployed on AWS, Azure, Google Cloud and mixed environments. This interoperability means that multi-cloud or hybrid organizations can implement the system without an architecture-level shock, and have uniform and verifiable auditing across all platforms.

VII. POTENTIAL APPLICATION

The proposed structure can be of paramount significance in the spheres in which the high compliance with the rules and auditability are critical. Some of the areas, which require round-the-clock attention to follow the standards, include finance, healthcare, and government which have to follow the standards, including PCI-DSS, HIPAA, GDPR, and ISO standards. The entities that work with autonomous auditors whose implementation is based on the use of the LLM (that are tied to the blockchain) would have an opportunity to have an uninterrupted and irreversible history of compliance operations and eliminate periodic manual audits. Such continuous monitoring reduces compliance drift, accelerates the certification process and ensures that the security controls are aligned with the regulatory requirements. At the same time, the framework improves the management of clouds by making autonomous discoveries of the misconfiguration of resources, identity abnormality, and unsecure deployment patterns. With the help of logs, telemetry, and analysis of access events contextual data, the auditors will be able to detect the threats in real-time and take or prescribe corrective actions with minimal human intervention.

Another powerful capability that is introduced by the system is the forensic investigation and incident responding. The multiple cloud telemetry can be correlated, and the root-cause story of convoluted paths of assaults can be created by the LLM auditors. These forensic outputs are then stored permanently in the blockchain creating a reliable and verifiable record of evidence that may be used in the regulatory investigations and the legal process. By incorporating autonomous reasoning, irreversible logging, and overall contextual analytics, the proposed solution have the potential to enhance post-incident learning and ensure the overall integrity of a forensic observation, which subsequently results in the overall efficiency of cloud security resilience.

VIII. DISCUSSION

The conceptualization of cloud security, implementation, and governance is a socio-technical change that is carried out by the merger of the autonomous auditors of LLM with the verification that is supported using blockchain. One of the key aspects is that auditors began to be more active unceasing learners in security, instead of passive assessors. In the traditional systems, an audit is conducted periodically and in a backward manner in which the vulnerabilities are only identified once they are used or when the misconfigurations have failed. The proposed architecture, in its

turn, will introduce the auditors into the life cycle of cloud systems, thereby allowing the identification of the evidence on-the-fly and immediate maintenance of the evidence. This transforms the time nature of auditing and in effect changes paradigm of episodic assurance by perpetual assurance.

Externalization of trust is the other important point of discussion. As opposed to centralized audit systems that rely on the trust to the cloud provider or the security teams of the company, the one supported by the blockchain permits the development of trust on the grounds of decentralized consensus and the cryptographic guarantee. It will reduce the use of third party representatives of the auditor and reduce the possibility of insiders influencing the audit results. This model offers a clear picture of security activities that is not subject to change by the regulators in industries that have a mandatory requirement to adhere to a certain regulation.

However, the connection between the auditors of LLM and blockchain infrastructure leads to architectural tensions, which should be handled with caution. The LLMs rely upon probabilistic reasoning and large context windows, and flexible interpretation of natural language policies, whereas blockchain systems rely upon strict determinism and unambiguous change of state. The solution to these radically distinct paradigms is to have controlled interfaces, e.g. smart contracts giving deterministic validation logic. This connection ensures that the invalidity of audit results is not adjusted to compliance states in the event of an LLM displays the hallucination or misjudgment of the layer of blockchain. In this regard, blockchain may be applied as confirmation anchor in which the cognitive output of an auditor is pegged on cryptographic assurance.

When it comes to the organizational environment, LLC auditors raise some grave governance and accountability concerns. Traditional audit teams rely on manual audit reviewers and human based auditor decisions can be subject to the laws and professional ethics. On the contrary, opaque decisions made by the neural network are applied by auditors of LLC. This would necessitate the introduction of AI control measures, like model elucidation interfaces, model hazard rating frameworks, training data provenance monitoring and model conduct confirmation customary. The completeness of the governance is also increased by the blockchain as the records of auditor operations are not alterable, the forensic audits can be replayed and the compliance checked, as well as the stakeholders involved.

The other point of discussion is the aspect of scalability. Large streams of telemetry are indeed generated by cloud environments and computational infeasibility of real-time LLM reasoning can be achieved. Model distillation, caching, and edge-cloud hybrid all inference methods are appropriate in the case of overhead reduction, although blockchain scalability ought to be considered, as well. Enterprise deployments should use permissioned blockchains because they are more throughput-intensive and can be easily configured to have controlled consensus mechanisms. On-chain hashing (Layer-2 anchoring) and off-chain storage can

be applied to address the bandwidth limitations and reduce the cost of operation.

Lastly, the social consequences of autonomous auditing systems should be considered. With the continuation of the tendency where organizations shift to AI-based compliance mechanisms, the auditing profession would be turned to supervising, where the human expert ensures coordination, verification, and regulation of autonomous agents instead of manual checks. This provides avenues of enhanced precision and efficiency and requires new moral standards, employee retraining programs, and regulation frameworks that consider machine-guided decision-making in the event of security assurance. On the whole, the discussion shows that the suggested LLM-blockchain convergence is not only a technological innovation but a strategic point of turning the history of cloud-security auditing.

IX. CONCLUSION

This publication introduces a thorough model of LLM-based autonomous auditors that are interconnected with blockchain-based infrastructures, which provides a prospective way of cloud security and compliance. The suggested system enables the conventional auditing methods to address the long-standing constraints of its predecessors such as centralized logs manipulation, late incident detection, and inconsistent reporting of compliance, by combining the interpretative intelligence of LLMs with immutability and verifiability of blockchain.

The long analysis has illustrated that the cognitive flexibility, the ability to think in context and to maintain continuous monitoring ensured by the LLM auditors are far much more than the rule based and manual auditing systems. Blockchain is used to strengthen these features with cryptographically secure and tamper-free trail of evidence and deterministic validation of audit events via smart-contracts. These technologies have the potential to provide a symbiotic ecosystem in which autonomous auditors can work with a high degree of assurance, transparency, and anti-manipulation.

In addition to this there is the alternative of introducing reinforcement learning, such that auditors will learn and get better with time as they are more accurate in detection, less false positive and remediation workflow. It also interoperates with the cross clouds, the privacy sensitive telemetry processing and regulatory alignment with it being the reason why the architecture can be applied in the industry where the main aspects are trust and verifiability.

Despite these encouraging advances, it is also acknowledged in the paper that there exist limitations in the form of computational overhead, scalability of blockchains, potential LLM hallucinations and governance. However, they do not deter the transformational abilities of the proposed model, but they signify the further domains of innovations and enhancement. The new directions are federated AI auditors, zero-knowledge compliance proofs, cryptographically bounded LLM reasoning and secure multi-agent interaction on multi-tenant cloud ecosystems.

In conclusion, self-sovereign blockchain-based auditors using LLM are an important advancement to the future of cloud security. They shift the paradigm of periodic, trust based audit to continuous, decentralized and verifiable assurance. This construct will provide a platform upon which intelligent auditors will collaborate with cloud systems to deliver real-time, impartial and hack-resistance security regulations that is relevant in the changing digital environment.

X. REFERENCES

1. S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008.
2. V. Buterin, "A Next-Generation Smart Contract and Decentralized Application Platform," Ethereum Whitepaper, 2014.
3. M. Swan, *Blockchain: Blueprint for a New Economy*, O'Reilly, 2015.
4. K. Christidis and M. Devetsikiotis, "Blockchains and Smart Contracts for the Internet of Things," *IEEE Access*, 2016.
5. Z. Zheng et al., "Blockchain Challenges and Opportunities," *Int'l Journal of Web and Grid Services*, 2018.
6. A. M. Antonopoulos, *Mastering Bitcoin*, O'Reilly, 2017.
7. D. Tapscott and A. Tapscott, *Blockchain Revolution*, 2016.
8. N. Szabo, "Smart Contracts," 1997.
9. C. Dwork and M. Naor, "Pricing via Processing or Combatting Junk Mail," *CRYPTO*, 1992.
10. ISO/IEC 27001, International Organization for Standardization, 2013.
11. NIST SP 800-53, U.S. Dept. of Commerce, 2013.
12. NIST Cybersecurity Framework, 2014.
13. A. Shamir, "How to Share a Secret," *Communications of the ACM*, 1979.
14. R. Rivest, A. Shamir, and L. Adleman, "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems," 1978.
15. D. Boneh, "Twenty Years of Cryptography," *Communications of the ACM*, 2008.
16. A. Juels and A. Opera, "Proofs of Retrievability," *CCS*, 2007.
17. L. Lamport, "Time, Clocks, and the Ordering of Events," *CACM*, 1978.
18. B. Schneier, *Applied Cryptography*, Wiley, 1996.
19. M. Al-Bassam, "Blockchain-Based Decentralized Cloud Computing," 2017.
20. J. Dean and S. Ghemawat, "MapReduce: Simplified Data Processing on Large Clusters," 2004.
21. J. MacQueen, "Classification and Multivariate Analysis," 1967.
22. Y. LeCun et al., "Deep Learning," *Nature*, 2015.
23. A. Vaswani et al., "Attention Is All You Need," 2017.
24. J. Devlin et al., "BERT: Pre-training of Deep Bidirectional Transformers," 2018.
25. I. Goodfellow, *Deep Learning*, MIT Press, 2016.
26. T. Mikolov et al., "Efficient Estimation of Word Representations," 2013.
27. R. Sutton and A. Barto, *Reinforcement Learning*, MIT Press, 1998.
28. S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall, 2009.
29. L. Bass, P. Clements, and R. Kazman, *Software Architecture in Practice*, 2012.
30. M. Armbrust et al., "A View of Cloud Computing," *CACM*, 2010.
31. A. Fox et al., "Above the Clouds: A Berkeley View of Cloud Computing," 2009.
32. E. Brewer, "CAP Twelve Years Later," *Computer*, 2012.
33. P. Mell and T. Grance, "The NIST Definition of Cloud Computing," 2011.
34. K. Hwang et al., "Cloud Security with Virtualization," *J. Security and Comm. Networks*, 2013.
35. S. Subashini and V. Kavitha, "A Survey on Cloud Security Issues," *Journal of Network and Computer Applications*, 2011.
36. M. Jensen et al., "Cloud Vulnerabilities," *IEEE ICSE Workshop*, 2009.
37. H. Takabi, J. Joshi, and G. Ahn, "Security and Privacy Challenges in Cloud," *IEEE Security & Privacy*, 2010.
38. J. Dastjerdi and R. Buyya, *Big Data Analytics: Distributed Systems*, Morgan Kaufmann, 2016.
39. R. Vadisetty, A. Polamarasetti, Aashna, S. K. Rongali, S. k. Prajapati and J. B. Butani, "Blockchain and Generative AI for Cloud Security: Ensuring Integrity and Transparency in Cloud Transactions," 2025 International Conference on Advancements in Smart, Secure and Intelligent Computing (ASSIC), Bhubaneswar, India, 2025, pp. 1-6, doi: 10.1109/ASSIC64892.2025.11158656
40. S. Chen et al., "Trust Management in Blockchain Systems," 2018.
41. X. Liang et al., "Integrating Blockchain with Cloud Computing," 2017.