International Journal of Computer Applications Technology and Research
Volume 15–Issue 09, 42 – 45, 2026, ISSN: -2319–8656
DOI:10.7753/IJCATR1509.1006

Sharing Fraud Intelligence Across Institutions: A Consortium Blockchain Framework for Banks and Mobile Money Providers in Kenya

Kevina BrenSda Mbati
Department of Computer Science and Information Technology
Kabarak University
Nakuru, Kenya

Ruth Oginga
Department of Computer Science and Information Technology
Kabarak University
Nakuru, Kenya

Nelson Masese
Department of Computer Science and Information Technology
Kabarak University
Nakuru, Kenya

Abstract: Fraud intelligence generated by one Kenyan bank or mobile money provider is rarely shared with peer institutions in a timely, trusted, and auditable way, leaving fraud schemes that move across institutional boundaries only partially visible to any single organisation. This study designs, implements, and evaluates a consortium blockchain framework that lets banks and mobile money providers submit, validate, share, and retrieve fraud intelligence within a permissioned, auditable environment. The framework is realized as a Hyperledger Fabric prototype spanning five simulated organisations, combining organisation-specific cryptographic identities, mutual TLS, chaincode-enforced access control, an append-only audit trail, and client-side hashing of account identifiers. Functional, security, and performance testing, using synthetic fraud-intelligence data and workloads of 50 and 100 transactions, show that the prototype supports the full fraud-alert lifecycle with a 100 percent transaction success rate, chaincode-enforced confidentiality, and a fully traceable audit trail, while transaction-volume performance improves rather than degrades over the tested range (mean response time falling from 390.5 ms to 304.0 ms). Adversarial testing confirms that identity-spoofing and unauthorized-access attempts are rejected, though a ledger-tampering test exposes a boundary condition in post-commit corruption detection. The main contribution is a working, empirically tested demonstration that permissioned consortium blockchain is a technically viable mechanism for controlled, auditable fraud-intelligence sharing among competing financial institutions, distinct from, and complementary to, institution-level fraud detection itself.

Keywords: Commercial banks; Consortium blockchain; Fraud intelligence sharing; Fraud prevention; Hyperledger Fabric; Kenya; Mobile money providers

International Journal of Computer Applications Technology and Research
Volume 15–Issue 09, 33 – 41, 2026, ISSN: -2319–8656
DOI:10.7753/IJCATR1509.1006

1. INTRODUCTION

Kenya's banks and mobile money providers operate an increasingly interconnected digital financial ecosystem, yet each institution's fraud detection remains confined to its own transaction data. When fraudsters move stolen funds across multiple banks and mobile money accounts, layering transactions to stay under any single institution's monitoring threshold, no institution sees the whole picture. Cyber fraud cases reported to the Central Bank of Kenya rose from 153 in 2023 to 353 in 2024, with losses rising from roughly KSh 412 million to KSh 1.5 billion over the same period [1], and SIM-swap and cross-institution fund-layering schemes were named as recurring, cross-boundary attack patterns during the same period [2]. The problem this study addresses is the absence of a secure, controlled, and auditable mechanism through which competing financial institutions can share fraud intelligence with one another.

This problem matters because the naive fix, simply asking institutions to exchange fraud data bilaterally or through a shared central database, runs into exactly the concerns that make institutions reluctant to share in the first place: none wants to hand a competitor, or a single central custodian, unrestricted access to its transaction records, and none can unilaterally guarantee that shared data will not be altered, misused, or exposed to unauthorized parties. A central database also concentrates risk: whoever controls it becomes a single point of failure and a single point of trust that every participating institution must accept. Left unresolved, this coordination failure means that fraud patterns detected by one institution routinely go unrecognized elsewhere, until the same scheme has already been repeated against several peer institutions.

Existing efforts have not solved this problem because they address adjacent parts of it rather than the coordination problem itself. Institution-level machine learning and rule-based detection [3] strengthens what a single institution can see in its own data but does nothing to extend that visibility across institutions. General blockchain applications in finance have concentrated on payment settlement and record-keeping rather than fraud-intelligence exchange, and where consortium blockchain has been proposed for threat-intelligence sharing in other sectors [4], it has not been adapted to, or empirically tested against, the specific governance and confidentiality constraints of banks and mobile money providers operating side by side in one regulatory jurisdiction. This study differs from both lines of work by treating fraud-intelligence sharing itself as the artefact to be designed, built, and tested, rather than treating blockchain as a payment rail or treating fraud detection as a purely institution-internal analytics problem.

The approach taken here is Design Science Research: literature and regulatory evidence were used to derive functional and non-functional requirements, which were then implemented as a permissioned Hyperledger Fabric network of five simulated organisations (Equity Bank, KCB, Co-operative Bank, Safaricom M-PESA, and Airtel Money), a fraud-intelligence chaincode enforcing submission, validation, sharing, and retrieval rules, a REST API gateway, and an analyst dashboard. The prototype was evaluated through functional, security, and performance testing rather than through claims alone. It demonstrated a complete, working fraud-alert lifecycle, confidentiality and transparency enforced at the chaincode level, and stable or improving performance from 50 to 100 transactions, while adversarial testing showed that spoofed identities and unauthorized queries were correctly rejected. The main limitations are that the underlying requirements were derived from literature rather than validated with practitioners, that only five organisations and two transaction workloads were tested, and that the study measures the technical feasibility of sharing fraud intelligence rather than a measured reduction in real-world fraud.

The remainder of the paper positions this work against related literature (Section 2), introduces a running example used throughout (Section 3), defines terminology (Section 4), presents the framework's design and implementation (Section 5), reports performance and security experiments (Section 6), and concludes (Section 7).

2. RELATED WORK

Hyperledger Fabric provides the permissioned, modular architecture on which most enterprise consortium blockchains, including this one, are built [5]; this study adopts, but does not modify, that underlying platform, and instead contributes an application-layer chaincode, gateway, and evaluation for a specific use case. TrustShare [4] demonstrates blockchain-based threat-intelligence sharing between organisations in a general security context; the present study adapts the same underlying idea, controlled multi-party sharing over a permissioned ledger, to the specific governance, confidentiality, and regulatory constraints of Kenyan banks and mobile money providers, and evaluates it empirically rather than only architecturally. Within Kenya specifically, prior work has examined fraud and privacy risk in mobile money from the perspective of individual agents and users [6] and has documented the scale of banking-sector cyber fraud [1], but neither proposes or tests a cross-institutional sharing mechanism. Finally, the wider blockchain-for-fraud-detection literature [3] cautions that blockchain by itself does not detect fraud; this study takes that caution as a design constraint rather than a criticism, and evaluates the framework explicitly as an intelligence-sharing mechanism rather than a fraud-detection algorithm.

3. RUNNING EXAMPLE

To keep the technical description concrete, the paper uses a single running example throughout: an analyst at Equity Bank detects a suspicious transaction linked to an account that has also transacted with an Airtel Money agent line, and wants to warn the other four consortium members without exposing the underlying customer data to anyone who has not been explicitly granted access. Sections 5 and 6 return repeatedly to this scenario, tracing how the alert is submitted, hashed, endorsed, ordered, committed, validated by a second institution, and then selectively shared, so that the same walk-through both introduces the framework's components and demonstrates their behavior.

4. PRELIMINARIES

This section defines terms used throughout the paper that are not original to this study. Design Science Research (DSR) is a research paradigm concerned with designing and evaluating artefacts that address an identified organizational problem [7]; this study follows the six-activity Design Science Research Methodology (DSRM) of Peffers et al. [8]: problem identification, definition of objectives, design and development, demonstration, evaluation, and communication. Hyperledger Fabric is a permissioned distributed ledger platform in which participating organisations are represented by a Membership Service Provider (MSP), which issues X.509 identities used to

authenticate every peer and client connection; transactions are executed by chaincode (smart contract logic), endorsed by peers according to a predefined endorsement policy, ordered into blocks by an ordering service, and committed to an append-only ledger shared across a defined communication channel [5]. The study evaluates the resulting artefact against five criteria: security (resistance to unauthorized access and tampering), confidentiality (restriction of intelligence to authorized parties), transparency (traceability and auditability of transactions), scalability (stability of performance as transaction volume, or consortium membership, increases), and fraud prevention effectiveness (the extent to which the artefact supports, rather than replaces, collaborative fraud response).

5. FRAMEWORK DESIGN AND IMPLEMENTATION

At the architectural level, the framework has three layers: a consortium blockchain layer, comprising five organisation-specific peers, a shared channel, and an ordering service; an application layer, comprising a Go-based chaincode and a REST API gateway; and a presentation layer, comprising an analyst dashboard used to submit, validate, share, and review fraud intelligence. Returning to the running example, when the Equity Bank analyst flags the suspicious account, the dashboard sends the alert, together with the account identifier, to the REST API gateway.

The gateway first hashes the sensitive account identifier using SHA-256, so that the raw account number is never written to the ledger, then assembles and signs a transaction proposal and submits it to the endorsing peers named in the network's endorsement policy. Each endorsing peer verifies the submitting organization's X.509 identity over a mutually authenticated TLS connection, executes the chaincode's SubmitFraudAlert function against its local copy of the ledger state, and returns a signed endorsement if the organisation is authorized and the submission is well-formed. The gateway forwards the endorsed transaction to the ordering service, which sequences it into a block; every peer on the channel then validates and commits that block, updating its local ledger and world state.

Once committed, the alert exists on every peer's ledger, but the chaincode's access-control logic still restricts who may read it: by default, a submitted alert is visible only to the submitting organisation. A second institution, say KCB, must independently corroborate or validate the alert through the ValidateFraudAlert function before Equity Bank can choose to invoke ShareFraudAlert, at which point the alert becomes visible to the organisations it is explicitly shared with. Every one of these actions, submission, validation, and sharing, is written as a new, timestamped, organisation-attributed ledger transaction rather than as an edit to an existing record, so that GetFraudAlertHistory can later reconstruct the complete lifecycle of the alert as an audit trail. Where multiple institutions corroborate the same account or pattern, the dashboard aggregates these independent submissions into a single elevated-confidence pattern, illustrating how the framework is intended to support, though not by itself perform, collaborative fraud detection.

6. PERFORMANCE AND SECURITY EXPERIMENTS

Performance experiments measured two things: pure running time, in the form of per-transaction response time and end-to-end cross-organisation dissemination latency, and sensitivity to an important parameter, namely transaction volume, by repeating the same workload at two sizes.

Two load-test batches of 50 and 100 SubmitFraudAlert transactions were each run three times; all six batches achieved a 100 percent transaction success rate. Table 1 reports response time and throughput at each size.

Table 1. Transaction Throughput and Response Time (N = 50 vs N = 100)

Metric	N = 50	N = 100
Success rate	100% (3/3)	100% (3/3)
Mean response time	390.5 ms	304.0 ms
Range (batch means)	301-528.7 ms	262-359.4 ms
Min / max (single tx)	166 / 2076 ms	162 / 2320 ms
Mean throughput	2.69 tx/s	3.31 tx/s

Mean response time fell from 390.5 ms at 50 transactions to 304.0 ms at 100 transactions, while mean throughput rose from 2.69 to 3.31 transactions per second, indicating that, over this range, increasing the workload did not degrade performance. Cross-organisation dissemination latency averaged 2680.9 ms across 50 share-and-confirm cycles (SD = 558.4 ms), with a median of 2523.0 ms and two upper-tail outliers (4168 ms, 5710 ms) attributable to occasional gossip-propagation delay rather than a systematic slowdown (Table 2).

Table 2. Cross-Organisation Dissemination Latency (N = 50)

Statistic	Value
Mean	2680.9 ms
Median	2523.0 ms
Minimum	2363 ms
Maximum	5710 ms
Standard deviation	558.4 ms

A second set of experiments tested sensitivity to adversarial conditions rather than to load. A transaction submitted with KCB's genuine credentials but a falsely declared EquityBankMSP identity was rejected during endorsement, confirming that Fabric validates the cryptographic identity of the submitter rather than a client-declared value. A query issued by SafaricomMSP for an alert it had not been granted access to returned a not-found response rather than the record itself, confirming that the chaincode's access control does not leak the existence of restricted records. A deliberate 16-byte corruption of one peer's local block file was not detected by ordinary read operations, a genuine limitation rather than a passed test, and is reported as such: it indicates a boundary condition in post-commit integrity checking that would need further investigation before the framework could be relied on for production-grade tamper detection.

Table 3. Evaluation Against the Five Criteria

Area	Rating	Interpretation
Security	4/5	Largely Satisfied
Confidentiality	4/5	Largely Satisfied
Transparency	5/5	Fully Satisfied
Scalability	3/5	Partially Satisfied
Fraud Prevention Effectiveness	4/5	Largely Satisfied

Table 4. Comparison of Current Practice and the Prototype

Criterion	Current Practice	Prototype
Security	Institution-specific systems with limited inter-organizational controls	Mutual TLS, X.509 identities, digital signatures

Confidentiality	Access depends on policy and bilateral agreement	Chaincode-enforced access control
Transparency	Limited shared trace of submission or sharing	Immutable, traceable transaction history
Scalability	Manual coordination harder as institutions grow	Volume scalability shown; institutional untested
Fraud Prevention Effectiveness	Intelligence remains fragmented	Validation, sharing, pattern support collaboration

What these experiments do not measure is institutional scalability, that is, performance as the number of participating organisations grows beyond five, since the development environment's available memory constrained testing to the five-organisation configuration used throughout the paper.

7. CONCLUSIONS

The experiments reported here make the paper's opening claim concrete: a permissioned consortium blockchain can carry fraud intelligence between competing Kenyan financial institutions with a 100 percent transaction success rate, a fully organisation-attributed and immutable audit trail, and response times that improve, rather than worsen, between 50 and 100 transactions, while correctly rejecting spoofed identities and unauthorized queries during adversarial testing. What the same experiments also make concrete is where the artefact's current evidence runs out: institutional scalability beyond five organisations, sustained availability under continuous load, and the post-commit ledger-tampering boundary condition are open questions rather than demonstrated properties, and no claim is made that the framework detects fraud or reduces fraud losses on its own. Read together, these results support a narrower and more defensible conclusion than "blockchain solves cross-institutional fraud": a permissioned consortium ledger is a technically workable substrate for the specific, previously unaddressed problem of sharing fraud intelligence across institutional boundaries in a controlled, auditable way, on top of which, rather than in place of, institution-level fraud detection can continue to operate.

8. ACKNOWLEDGMENTS

I take this opportunity to thank Dr. Ruth Oginga and Dr. Nelson Masese, Department of Computer Science and Information Technology, Kabarak University, for their supervision and guidance throughout the design, development, and evaluation of this work, also Dr.Andrew Kipkebut and Dr.Rose Sagwe for reviewer feedback on an earlier draft of the underlying thesis. This manuscript is derived from a Master of Science thesis undertaken at Kabarak University; no external funding was received.

9. REFERENCES

[1] Central Bank of Kenya. (2025). Kenya financial sector stability report 2024. Central Bank of Kenya.

[2] INTERPOL. (2026). African cyberthreat assessment report 2026. INTERPOL.

[3] Farrukh, H., Zafar, S., Rehman, Z. U., Shah, A. A., and Alshammry, N. (2025). Blockchain-based fraud detection: A comparative systematic literature review of federated learning and machine learning approaches. Electronics, 14(24), Article 4952.

[4] Ali, H., Buchanan, W. J., Ahmad, J., Abubakar, M., Khan, M. S., and Wadhaj, I. (2025). TrustShare: Secure and trusted blockchain framework for threat intelligence sharing. Future Internet, 17(7), Article 289.

[5] Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., Manevich, Y., Muralidharan, S., Murthy, C., Nguyen, B., Sethi, M., Singh, G., Smith, K., Sorniotti, A., Stathakopoulou, C., Vukolic, M., et al. (2018). Hyperledger Fabric: A distributed operating system for permissioned blockchains. Proceedings of the Thirteenth EuroSys Conference. https://doi.org/10.1145/3190508.3190538

[6] Sowon, K., Munyendo, C. W., Klucinec, L., Maingi, E., Suleh, G., Cranor, L. F., Fanti, G., Tucker, C., and Gueye, A. (2024). Design and evaluation of privacy-preserving protocols for agent-facilitated mobile money services in Kenya. Computers and Security, 133, Article 103569. https://doi.org/10.1016/j.cose.2023.103569

[7] Hevner, A. R., March, S. T., Park, J., and Ram, S. (2004). Design science in information systems research. MIS Quarterly, 28(1), 75-105. https://doi.org/10.2307/25148625

[8] Peffers, K., Tuunanen, T., Rothenberger, M. A., and Chatterjee, S. (2007). A design science research methodology for information systems research. Journal of Management Information Systems, 24(3), 45-77. https://doi.org/10.2753/MIS0742-1222240302