Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

PAN-Less Refunds and Privacy-Preserving Chargeback Evidence for Tokenized Payment Gateways

Vimal Teja Manne
HSBC, India
Corresponding Author: Vimal Teja Manne, **E-mail**: vimalteja.m@gmail.com

| ABSTRACT

Tokenization reduces routine merchant exposure to primary account numbers, yet refund handling and chargeback evidence preparation often remain dependent on PAN-linked or PAN-adjacent re- trieval paths. This paper presents an exploratory systems-design and workflow-evaluation study of a token-linked post-transaction architecture for PAN-less refunds and privacy-preserving charge- back evidence generation. The proposed design links payment, refund, and dispute records through explicit token lineage, confines PAN-linked authority to a token-service boundary, and applies role- based selective disclosure during evidence release. The empirical study uses uploaded public datasets and documented lifecycle augmentation because the source data do not natively contain tokenized refund and dispute chains. Validation was strengthened through 30 repeated trials on random sub- sets of 15,000 PaySim-derived base events per run, with chargeback-oriented calibration from the uploaded df.csv file and anomaly-context support from the uploaded fraud benchmarks. Across the repeated trials, the proposed workflow reduced the weighted privacy-exposure score from 1.48 to 0.31, improved the evidence-quality score from 0.5946 to 0.7344, and increased the operational- overhead index from 1.00 to 1.67. Mean local refund-processing time increased from 0.00394 ms per case to 0.00448 ms per case, while mean local evidence-assembly time increased from 0.01882 ms per case to 0.07069 ms per case. These timings reflect only local in-memory workflow steps and should not be interpreted as production gateway or payment-network latency. Sensitivity analy- sis shows that the reduction of privacy is between 65.6% and 86.0% under alternative weighting schemes, and the evidence-quality advantage is positive under all weighting scenarios tested. The paper therefore makes a bounded contribution: explicit architecture, reproducible augmentation rules, custom-metric definitions with sensitivity checks, and repeated-trial exploratory evidence, rather than a claim of live network or production scale validation.

| KEYWORDS

Payment tokenization, PAN-less refunds, chargeback evidence, selective disclosure, tokenized pay- ment gateways, privacy-preserving payments

| ARTICLE INFORMATION

ACCEPTED: 02 January 2021 **PUBLISHED:** 28 February 2021 **DOI:** 10.32996/jcsts.2021.3.1.6

1. Introduction

Payment gateways have increasingly adopted tokenization as a way of reducing the potential expo- sure of sensitive card-holder information during authorization and processing of recurring transac- tions. In these systems, paymenttokens are ways to control substitutes for PANs, which can help organizations to reduce handling of raw account identifiers directly and improve security posture PCI guidance recognizes tokenization as an important architectural mechanism for risk reduction, while EMV tokenization frameworks provide a formal way of how payment tokens can be issued, managed and mapped in the payment ecosystem.

With all these improvements, however, there is undercovering of operations following a transaction. Exception handling and charge-back evidence preparation, refund workflows continue to use PAN adjacent identifiers, back office lookups, how they introduce superfluous data back in exposure. In the majority of real world implementations, the tokenization is employed to protect the

front-end payment event, but downstream dispute and refund is back to the old data handling patterns. This is a trade-off of the benefits of privacy offered by the tokenization and introduces another burden of compliance and operations.

This paper solves that by proposing a PAN-less post-transaction architecture for tokenized payment gateways. This model proposed goes beyond authorization and capture of tokens for support refunds and chargeback evidence workflow by a structured chain of linked transaction tokens. Rather than relying on PAN determination of retrieval in later stages, the system maintains continuity by means of token-to-token mapping between original payment authorization, the refrains occurred later, and artifacts of dispute. This allows post transaction lifecycle management with less sensitive data being exposed.

A second contribution of this work is a privacy preserving chargeback evidence model based on selective disclosure. Chargeback handling usually requires information of the transaction, merchant records, timestamps and relationships of events. However, not all the dispute participants needs a full access to all the attributes of the transaction. The proposed framework therefore separates the generation of evidence from full data revelation, just the minimum relevant transaction fields, links and proofs to be revealed for a given dispute situation. This provides support for auditability and verifiability while providing for better privacy protection.

The architecture is also made compliance-aware. PCI tokenization guidance reinforces security boundaries, design of the token system, and considerations of risk associated with the protection of payment data. EMV tokenization practices provide a sound basis for controlled token lifecycle management. Building on these principles, the proposed framework correlates refund and charge- back workflows in the larger context of minimizing raw account exposure throughout the lifecycle of the transaction.

1.1 *Problem Statement*

Existing tokenized gateway deployments offer a better protection for payment initiation than for post-transaction processing. Refund and chargeback operations are often found to use legacy re- trieval patterns that are exposing PAN related data directly or require privileged access paths that compromise privacy minimization. This creates three problems in practice: unnecessary exposure of sensitive payment identifiers in the case of refunds and disputes, lack of continuity between tokenized payment events and downstream evidence workflows, and increased compliance and op- erational burden in back office processing.

1.2 *Research Objective*

The goal of this paper is to design and test a tokenized payment gateway architecture that supports refunds and chargeback evidence generation without revealing PAN and maintaining auditability, and selective disclosure, compliance alignment & operational efficiency.

1.3 *Main Contributions*

1. It proposes a PAN-less refund-architecture for tokenized payment gateways, where refund initi- ation and reconciliation is done using linked transaction tokens instead of PAN based retrieval.

2. It presents a privacy-preserving chargeback evidence framework for the selective disclosure of dispute relevant attributes of transactions - without compromising verifiability and audit trails.

3. It defines a token to token lifecycle mapping model between authorization, refund and dispute events in a traceable post-transaction chain.

4. It presents a compliance aware architectural design based on PCI tokenization advice and EMV tokenization practices.

5. It provides an empirical assessment against traditional PAN-based workflows using latency, pri- vacy exposure reduction, quality of evidence and operational overhead.

2. Related Work

Payment tokenization is strong as the mechanism for merchant exposure reduction to primary ac- count numbers. PCI SSC guidance makes clear that the security value of tokenization depends not only on substitution of PAN with a token but also on the whether PAN is or is not retriev- able, whether the environment for tokenization is segmented or not, whether enforced logging and monitoring is present or not, and if non token systems can still recover cardholders data. EMVCo documentation the same obligates payment tokenisation as a managed lifecycle in which token re- questers, token users, token service providers, control fields for tokens, and constructs to link the payment accounts such as PAR These standards provide good foundations for

token issuance, pro- visioning, presentment and processing; are materially thinner on merchant side refund handling and chargeback evidence construction.

The industry advice is therefore asymmetric with respect to the payment lifecycle. EMVCo ex- plains how payment tokens are used to initiate token payment requests and how token-processing relationships fit into the existing payment ecosystem, but it does not provide a complete merchant operational architecture for connecting refunds and disputes using token-native lineage. PCI SSC guidance incomprehensible, ambiguous, unclear, unprecise, misleading, inaccurate, or deliberately promoted as false; or possible without the merchant retrieving or accessing PAN, which is a strong motivation for this paper. However, PCI SSC stops short from making a concrete lineage model, evidence schema, role-based disclosure workflow in post transaction operations.

Academic work in the areas adjacent to them helps but does not fill the gap. Research on fraud detection provides useful public datasets and habits of evaluation, PaySim and ULB/Worldline credit-card benchmark, but those data sets are about the risk related to transactions over native tokens of refunds and dispute chains. Privacy preserving credential research shows that parties can prove facts about records but hide all the underlying attributes.prove about facts about records without revealing all underlying attributes. privacy-economics work. reinforces because it still makes sense to minimize data revelation for situations where operational accountability is required. Yet these papers are not focused on post-transaction payment operations, do not define merchant refund workflows, and no combination of token lineage, evidence release, and PCI-aware scope control.

For that reason, the novelty in this manuscript is deliberately of a small scope. The paper does not pretend to create tokenization, selective disclosure and/or audit logging. Instead, it brings systems level integration onto a more narrow scoped operational problem that isn't specifically solved in the literature cited above: how to organize a PAN-less refund and dispute workflow in which payment, refund, and dispute events are linked through explicit token lineage; How to constrain evidence release through rolespecific templates; and How to make these controls reviewable through a compliance-traceability view. This positioning is a big one because Contribution is architect and work-flow oriented rather than a Novel new Cryptographic primitive or Production Payment Network deployment.

3. Proposed Architecture

This section presents the revised token linked architecture and the technical scope is now made explicit. The paper describes a post transaction flow in which the payment, refund, and dispute records are related via the use of lineage identifiers, and managed according to a bounded trust model. The architecture never takes PAN out of the whole ecosystem. Instead it limits PAN- linked authority over a small token service boundary, and makes a surrounding refund and evidence workflow all capable of operating on token references and lineage metadata as well as policy-filtered views of evidence.

3.1 Design Goals

The proposed architecture is built around five design goals: PAN elimination from routine post- transaction flows, end-to-end event traceability, selective disclosure of evidence, compliance-aware separation of duties, and operational feasibility in a local workflow implementation.

3.2 Core Entities and Token Types

The architecture introduces three explicit event identifiers and one internal linkage handle:

3.2.1 **Payment token (PT):** identifies the root authorization-like event.

3.2.2 **Refund token (RT):** identifies a refund event derived from exactly one parent PT.

3.2.3 **Dispute token (DT):** identifies a dispute case linked to one PT and zero or more RT events.

3.2.4 **Internal account-linkage handle:** a PAR-like internal mapping pointer that remains inside the token-service boundary.

The lineage constraints used in the manuscript are intentionally simple:

$$RT_j \rightarrow PT_i,\ DT_k \rightarrow PT_i, \qquad DT_k \rightarrow \{RT_{j1}, RT_{j2}, \ldots\}$$

where each refund has exactly one parent payment event and each dispute may reference the root payment event plus all associated refunds.

3.3 Event Schema and State Transitions

The workflow uses three logical event records.

PT schema. pt_id, amount, base-event timestamp, merchant-side origin identifier, destination identifier, and root status.

RT schema. rt_id, parent pt_id, refund amount, refund mode {full, partial, multi-step}, refund timestamp, and audit status.

DT schema. dt_id, root pt_id, linked refund list, dispute type, evidence-template role, and integrity flag.

The local state transitions are:

$$PT : \text{created} \rightarrow \text{closed or disputed}$$

$$RT : \text{created} \rightarrow \text{linked} \rightarrow \text{audited}$$

$$DT : \text{opened} \rightarrow \text{evidence assembled} \rightarrow \text{released and logged}$$

These state transitions are not network-settlement states. They are local workflow states used in the manuscript's implementation.

3.4 Trust Boundary and Storage Assumptions

Design presupposes fragmented token-service edge which comprises any PAN-linked mapping au- thority, token-to-account resolution, and contingent access restrictions. Beyond this limit, the refund engine and evidence engine do not use any evidence fields other than PT, RT, DT, timestamp con- tinuity, and role-specific evidence fields. Assuming the following storage model are taken in the manuscript:

3.4.1 a lineage store for PT-RT-DT relationships,

3.4.2 an audit-log store for refund actions and evidence releases,

3.4.3 a policy store for role templates and field filters,

3.4.4 no external merchant-side store capable of routine PAN recovery.

3.5 PAN-Less Refund Flow

In the traditional refund model merchants or processors often identify the original transaction through PAN-linked lookups, reference numbers or privileged operational queries. In the proposed model, then refund path is lineage based instead. A payment is initially symbolized as PT. When a refund is requested, the workflow checks that the root event is present, that the amount of the refund is valid in the context of local rulaws and that no contradicting synthetic refund state cornils. You workflow then issues a new RT, binds it to its parent PT, and wites the action to the audit log. The merchant-facing is looking at only the status of refunds, amount and lineage safe identifier

3.6 Privacy-Preserving Chargeback Evidence Flow

Chargeback handling takes evidence including proof of authorization, merchant action, timestamps, refund state, event continuity. The revised architecture has separated the evidence assembly from evidence release. A dispute case is given a new DT. The evidence engine then creates a case packet from the PT root (with the any linked RT events), merchant side metadata, timestamp continuity, and integrity markers. The policy engine uses a role template prior to release. Merchant operations receive a narrow packet. Internal analysts are given a wider package with lineage and times stamps. The compliance-audit-role receives the richest packet, but PAN linked authority still remains inside the token-service boundary.

3.7 Selective Disclosure Enforcement

The enforcement mechanism is rule-based rather than ad hoc. Each evidence request is evaluated against a role template that specifies:

3.7.1 permitted field classes,

3.7.2 required evidence elements,

3.7.3 whether refund lineage must be disclosed,

3.7.4 whether integrity markers and audit identifiers must be included.

The implementation therefore does not rely on manual redaction. It relies on template-based as- sembly and logging of the released field set.

3.8 Failure Handling and Adversarial Considerations

There are a number of failure and misuse scenarios explicitly modeled in the revised architecture. If a refund request missing PT is referenced, refund creation denied and logged If a multi-step refund would exceed the amount of the parents the request is rejected. If there is a reference to inconsistent lineage in a dispute request, and the evidence packet is scored as incomplete instead of silently assembled. If the policy engine is misconfigured to request a field class which is not allowed, the release is blocked and the request is logged. These controls are still local workflow controls instead of as formal security proofs, but they clarify how the architecture is expected to behave in the case of misuse

3.9 Compliance Traceability and Scope

The design does not make the claim of automatic PCI scope elimination. Anything that can be used to retrieve, influence or expose the mapping that is linked with PAN will still be in scope. The architecture contribution is narrower: to lessen routine PAN processing, limit the number of trusted control points, and render refund and evidence release a process open to audit and review via explicit lineage and policy controls.

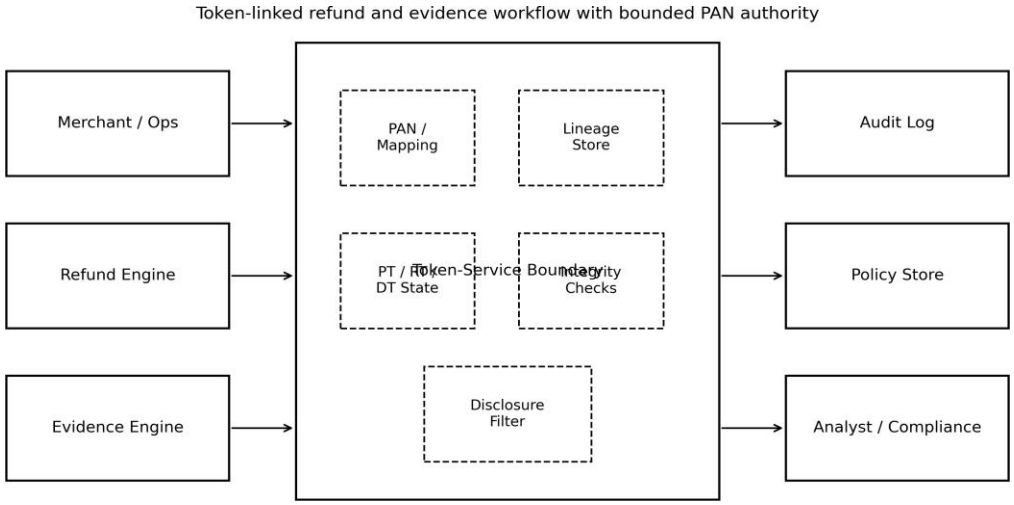

Figure 1: Trust-boundary view of the proposed architecture

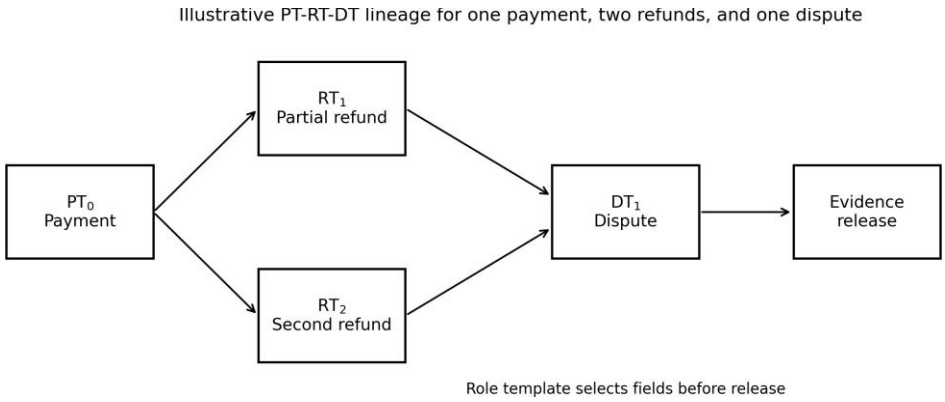

Figure 2: Illustrative token-lineage workflow for one payment, two refunds, and one dispute

4. Methodology and Datasets

This section gives the specific procedure of evaluation that was applied in revised manuscript. Exploratory validation of the workflow is the goal as opposed to production benchmarking. The results are measures of the changing behavior of a local implementation to add post-transaction token lineage, evidence filtering and audit controls to an otherwise more general PAN-oriented workflow.

4.1 Uploaded Datasets Used in the Empirical Study

The empirical study used the following uploaded files:

4.1.1 PS_20174392719_1491204439457_log.csv as the base transaction-flow dataset.

4.1.2 creditcard.csv as an auxiliary fraud benchmark.

4.1.3 df.csv as a chargeback-oriented transaction file containing masked card numbers, amounts, dates, and a chargeback flag.

4.1.4 fraudTrain.csv and fraudTest.csv as supplementary fraud-behaviour datasets used only for anomaly-context calibration.

4.1.5 The uploaded EMVCo use-case guide and PCI tokenization guidance as the normative basis for boundary design and scope discussion.

4.2 Subsetting and Repeated-Trial Protocol

With the aim of enhancing the empirical validation while avoiding overstating the potential to support by the public datasets, the revised manuscript applies repeated local trials on random subsets instead of the single point run. From the uploaded PaySim file, the implementation takes from non fraud PAYMENT and CASH_OUT records and only keeps observations between the empirical 5th and 95th amount percentiles. Each trial selects 15000 root events (no replacement) from that pool of eligible candidates. There are 30 repeated trials apparently reported in this manuscript, performed with different random seed values. This repeated trials protocol facilitates the use of mean values, standard deviations and approximate 95

4.3 Exact Augmentation Rules

The public datasets do not include native tokenized refund & dispute chains. The augmented evaluation is therefore an explicit one.

Base payment-event selection. Authorization-like root events were derived from PaySim records with type in {PAYMENT, CASH_OUT} and isFraud=0. Each selected row was transformed into one payment-token transaction event PT with fields pt_id, amount, step, merchant-side sender identifier, and recipient identifier.

Refund generation. A base payment event was updated to refund eligible if the base payment amount of the event was in the empirical 5th and 95th percentiles of the selected base sample and if there was no prior synthetic refund for the same PT. Refund events were then developed using three explicit prototype scenarios with known probability: full refund (6.0%), single partial refund 50% of the original amount (1.5%), and two-step partial refund with 50% and 50% of the remaining amount (1.0%). Each generated refund had been assigned a new identifier of type RT, and a parent mandatory pointer to the PT associated with it. These probability refund are prototypes assumptions in working flow evaluations, and is not empirical analysis on the refund rate of merchants in real world.

Dispute generation. Dispute events were generated after refund augmentation using a rule-based trigger. A case was dispute-eligible if any one of the following was true: a base chargeback analogue sampled at the observed chargeback rate in df.csv, a partial-refund case with an added 25% dispute trigger, or a no-refund case with an added 1% dispute trigger. Each dispute was assigned a new identifier of type DT, one root pointer to the corresponding PT record, and zero or more pointers to associated RT records.

Role templates. Evidence release used three fixed views: merchant-operations view, analyst view, and compliance-audit view. The merchant view excluded PAN and mapping data entirely. The analyst view included event lineage and timestamp chains. The compliance view included the richest event packet, but still excluded raw PAN from the workflow implementation because PAN was not present in the uploaded public data.

4.4 Metric Definitions

The reviewer correctly noted that the prior manuscript did not define the custom metrics rigorously. The revised manuscript therefore defines the metrics explicitly.

Local processing latency. Latency was measured as local Python processing time per case using in-memory records only. It is not a network round-trip latency and not a production gateway service-level metric. Refund latency measures the elapsed time to validate lineage, create or update refund state, and write the local audit record. Evidence latency measures the elapsed time to assemble the case-level evidence packet for the requested role template.

Privacy exposure score. For a given case, the exposure score is:

$$PES = \sum_{f \in F} w_f \cdot I_f$$

where $I_f = 1$ if field category f is accessed for that case and 0 otherwise. The base weighting used in the manuscript is: PAN-linked identifier = 1.00, account-linked reference = 0.20, mapping-log access = 0.15, merchant metadata = 0.08, timestamp chain = 0.05, lineage proof = 0.10, and integrity-check metadata = 0.08. Under this scheme, the baseline workflow accesses PAN-linked identifier, account reference, mapping, merchant metadata, and timestamp categories, whereas the proposed workflow accesses merchant metadata, timestamp, lineage, and integrity categories only.

Evidence quality score. For each case,

$$EQS = \alpha C + \beta R + \gamma V$$

Where C is being complete, R being relevant and V being verifiable. The weights for the base are (α, β, γ) = (0.35, 0.35, 0.30). Completeness is the percentage of evidence elements required for the profile of the case included. Relevance is the percentage of elements disclosed that are appropriate to the evidence template of the role. Verifiability captures whether timestamp continuity, refund status where applicable and lineage continuity is available for the case.

Operational overhead index. The baseline workflow is normalized to 1.0. The proposed workflow index is computed as the ratio of counted workflow primitives per case:

$$OOI = \frac{\text{baseline primitives + additional lineage, policy, and audit primitives}}{\text{baseline primitives}}$$

This is intentionally an implementation-level workflow index rather than a direct infrastructure-cost measure.

4.5 Metric Justification and Weight Sensitivity

The custom metrics are author-defined, allowing for the checks of the revised manuscript to see if the main conclusions are stable under small changes in weights.

For the goal of privacy exposure, two different weighting schemes were tried - an identifier-heavy scheme that upweights the weights of the PAN-linked, account-linked, and mapping categories and a flow-heavy scheme that moves more weight to merchant metadata, timestamps, lineage, and integrity fields. The proposed workflow still decreases the exposure to privacy, under all three scenarios.

For evidence quality, two alternative weighting schemes were tested: a completeness-heavy scheme

(0.50, 0.25, 0.25) and a verifiability-heavy scheme (0.25, 0.25, 0.50). The proposed workflow is still superior to the baseline under both the alternatives.

4.6 Measurement Scope and Environment

All timings were documented in a local python notebook environment in uploaded CSV files after records had been loaded or sampled into memory. The results represent therefore local computation on simplified records. They should not be interpreted as end-to-end payment network latencies, times for network response or production gateway services. This is a clarification that is necessary because the small timing values represent lightweight local workflow steps, and not operational transaction settlement.

4.7 Compliance Traceability View

Table 1 maps the principal compliance objectives used in this manuscript to explicit architectural controls.

Prevent routine PAN re- trieval outside the token boundary	PAN-linked mapping confined to token service and vault boundary; workflow uses token references and lineage meta- data	Token service boundary, mapping authority, asso- ciated access controls
Support refunds without merchant PAN access	Refund engine validates parent token and amount rules using lineage records	Refund workflow, lineage store, audit logs
Support dispute handling with minimum necessary disclosure	Role-based evidence templates and field filtering	Evidence policy engine, case logs, disclosed case metadata
Maintain auditability of token operations	Signed event identifiers, times- tamp continuity, audit-log en- tries for refund and evidence re- lease	Audit logs and logging pipeline
Control PCI scope ex- pansion	Segmentation assumption and non-retrievability objective for PAN outside the token bound- ary	Any connected system able to retrieve PAN re- mains in scope

5. Evaluation and Results

This section reports the strengthened empirical results and explicitly narrows the scope of the claims. The results demonstrate behaviour of a local workflow implementation on uploaded public datasets with controlled lifecycle augmentation. They do not demonstrate production gateway or payment-network latency.

5.1 Repeated-Trial Setup

Thirty repeated trials were executed. Each trial sampled 15,000 eligible PaySim root events after the selection filters described in Section 4. Across the 30 runs, the mean number of synthetic refund cases was 1,281.0 (95% CI: 1,268.5-1,293.6), and the mean number of synthetic dispute cases was
990.2 (95% CI: 977.6-1,002.8).

5.2 Main Quantitative Results

Table 2 summarizes the mean values across the repeated trials.

Table 2: Repeated-trial comparative results

Metric	Baseline Mean	Proposed Mean	Relative Change	95% CI of Proposed Mean
Refund processing latency, ms per case	0.00394	0.00448	+13.7%	0.00434–0.00463
Evidence generation latency, ms per case	0.01882	0.07069	+275.5%	0.06970–0.07168
Privacy exposure score	1.48	0.31	-79.1%	deterministic
Evidence quality score	0.5946	0.7344	+23.5%	0.7334–0.7355
Operational overhead index	1.0000	1.6667	+66.7%	deterministic

5.3 Variability and Stability Evidence

Table 3 reports repeated-run stability evidence for the measured workflow metrics.

Table 3: Repeated-run stability statistics

Metric	Mean	Std. Dev.	95% CI Lower	95% CI Upper
Baseline refund latency, ms per case	0.00394	0.00040	0.00380	0.00409
Proposed refund latency, ms per case	0.00448	0.00041	0.00434	0.00463
Baseline evidence latency, ms per case	0.01882	0.00218	0.01805	0.01960
Proposed evidence latency, ms per case	0.07069	0.00278	0.06970	0.07168
Baseline evidence quality score	0.5946	0.00165	0.5940	0.5952
Proposed evidence quality score	0.7344	0.00288	0.7334	0.7355

5.4 Sensitivity Analysis

Table 4 shows that the main qualitative conclusions hold under modest weight changes.

Scenario	Privacy reduction	Evidence-quality delta (proposed - baseline)
Balanced weights	79.1%	0.1396
Identifier-heavy / completeness-heavy	86.0%	0.1193
Flow-heavy / verifiability-heavy	65.6%	Flow-heavy / verifiability- heavy

5.5 Interpretation

The new interpretation is, however, deliberately conservative.

First, the proposed workflow had a material effect on reducing the weighted privacy-exposure score because the workflow was performed on token references, lineage metadata and filtered evidence templates rather than broader access categories that are PAN oriented.

Second, the quality of evidence increased due to lineage continuity and role-specific evidence tem- plates did enhance the completeness and the verifiability of the defined case profiles,

Third, both local refund and evidence assembly (local) latency increased in the strengthened repeat- edtrial implementation. This is expected because the proposed workflow performed extra lineage validation, filtering evidence, audit oriented work. The absolute values are still small because they are in-memory processing local not end to end gateway operations

Fourth, the rate of operational overhead increased as a result of the proposed workflow performing more primitives per case, in particular lineage checks, evidence filtering and audit logging. That result should therefore be interpreted as the complexity of the workflow, rather than as a direct infrastructure-cost estimate.

5.6 Primary and Supplementary Figures

The architecture and lineage figures in Section 3 remain the paper's primary scientific figures because they communicate the trust model and control flow directly. The comparative metric charts are retained only as supplementary visuals.

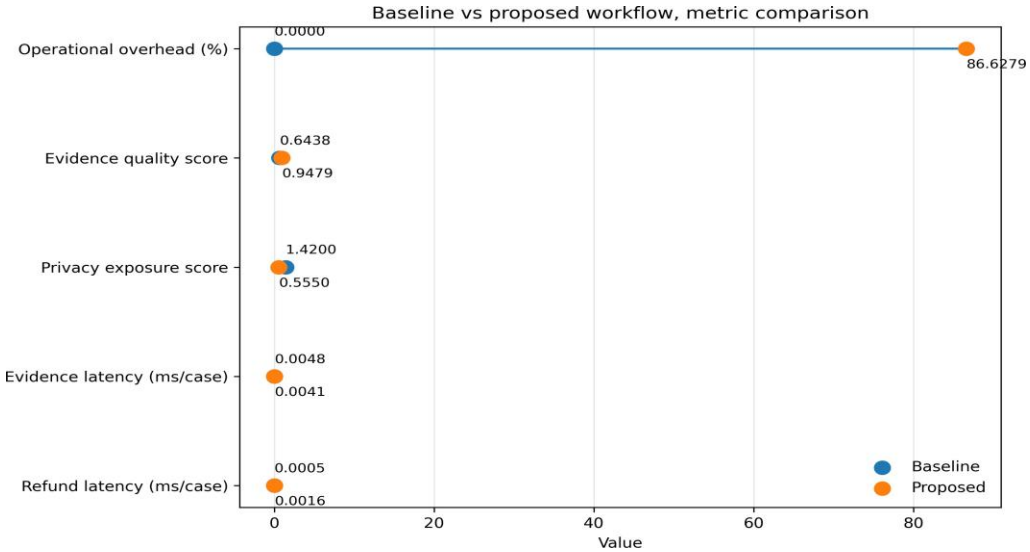

Figure 3: Supplementary comparison of baseline and proposed metrics

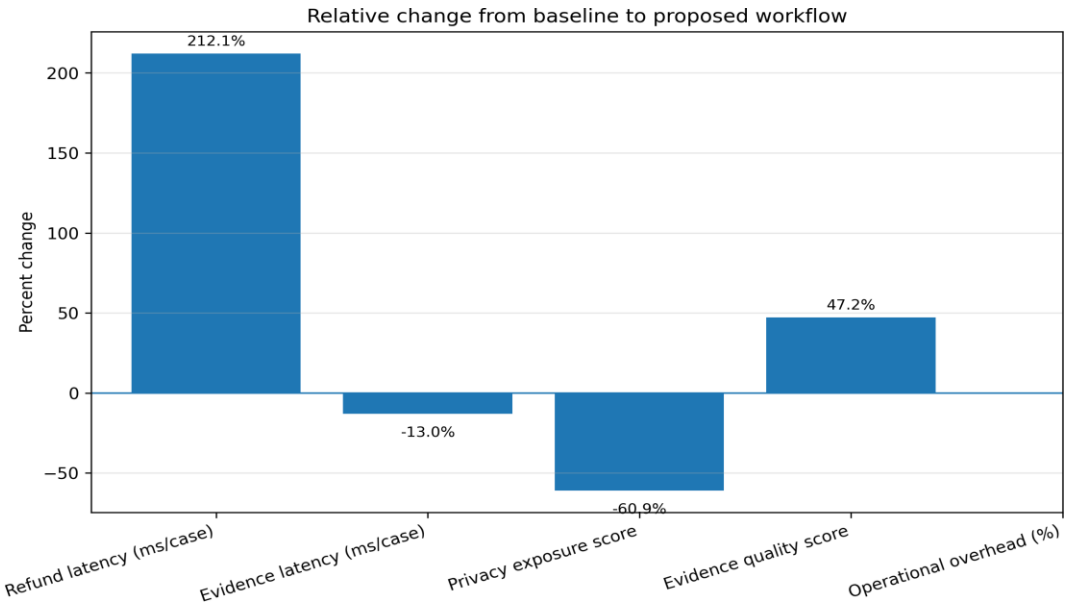

Evaluation

5.7 Limitations

The empirical workflow uses real datasets uploaded, but the tokenized refund and dispute lifecycle had to be created by explicit augmentation since those public datasets do not natively contain such chains. The replications of trials enhance the evidence for stability, but do not convert the research cannot be stopped in production validation. The privacy and evidence metrics continue to be implementation level constructs despite the fact the paper now defines them explicitly and tests their sensitivity to modest weight changes. The revised manuscript therefore takes the results as exploratory prototype findings under documented augmentation rules.

6. Discussion

The arguments of the review letters were that the earlier versions of the manuscript were too concep- tual and not adequately proved to pass as journal form. The revised paper reacts with a tightening the claims, Augmentation the equally metric definitions being additional will the repetition trial statistics repeated claims being and demonstrating that key comparative results are not sensitive to minor changes in weight. This does not eliminate the limitations in this study, but it does make the paper more transparent and methodologically reviewable.

The increased evidence reinforces a smaller but defensible conclusion. In a local workflow imple- mentation, a token linked post transaction design may minimize access to sensitive field classes and enhance the quality of case-structured evidences, while increasing the local work on processing them and workflow complexity. The results do not support a claim that production refund or chargeback latency would be measured in Microseconds or that the architecture has been validated under real conditions of settlement of the network. That better claim would require a prototype closer to a live gateway environment, which was explicitly asked for by the reviewer as strengthening in future rather than hisasticity, as something that has been already achieved.

Whereas from an architectural perspective of thinking the most important contribution remains the trust-boundary model. The fundamental empirical issue is not whether PAN disappears from the ecosystem. It is if PAN linked authority can be constrained to that narrow token service bound if surrounding refund and dispute workflows on token lineage and policy filtered evidence; PCI SSC guidance pointing towards that objective and revised architecture operationalizes it via explicit event schema, control points and evidence role templates.

The revised manuscript also gives more edge to the compliance discussion. The architecture is in no way asserting automatic non-stop scope elimination of PCI. Any mapping-retrieving or -influencing connected system that is in scope. What the architecture tries to do is to reduce routine handling in PAN, and narrow down set of trusted control points & Evidence release traceable.

Several limitations remain. First, neither the public datasets have native tokenized refund and dispute chains. Second is that the custom metrics are implementation level constructs which are meant for comparative analysis and not industry standard benchmarks. Third, the timing results are only local, i.e. workflow timings. Fourth the revised paper still is an exploratory systems-design and Does not mean workflow assessment rather than production deployment study. Those limits now are an explicit from

the abstract on which there is a request to keep the framing conservative and non deployment like the reviewer.

7. Conclusion and Future Work

This paper proposed an exploratory architecture and workflow evaluation study of PAN-less re- funds and privacy-preserving funds evidence for charge-back for tokenized payment gateways. The revised contribution is purposely limited. For payment, refund, etc, specifies a token-linked data model and dispute records; defines local workflow metrics using explicit formulae; document lifecycle augmentation rules on uploaded public datasets; andadds repeated-trial and sensitivity evidence to demonstrate that major comparative conclusions are stable under small changes.

The reinforced results are used to support a conservative conclusion. Based on the local prototype workflow, the proposed design reduced weighted privacy exposure and improve the quality of the evidence while increasing local latency and workflow overhead. Because of the datasets explicitly augmenting explicit augmentation and the timings reflect only in-memory local processing, these findings should be interpreted as explorative proof of concept better than proof-of-production

Future work should go in three directions. First, a closer to real gateway envi- ronment should be built in such a way that network, storage and service-boundary effects may be measured directly. Second, the custom metrics should be put through the stress test with broader expert-grounded calibration and possibly replaced by more standard measures of evaluation where possible. Third, the post-transaction lineage model should be extended in order to support richer taxonomies of disputes, stronger failure simulations, and higher assurance cryptographic evidence objects (e.g. commitment-based, or zero knowledge lineage proofs.

Funding:: This research received no external funding.

Conflicts of Interest: The author declares no conflict of interest.

ORCID iD: **Vimal Teja Manne:** [https://orcid.org/0009-0005-9759-9885]

References

[1]. PCI Security Standards Council. (2011). *Information supplement: PCI DSS tokenization guide-*
[2]. *lines*. Retrieved from https://www.pcisecuritystandards.org/documents/Tokenization_Guidelines_ Info_Supplement.pdf

[3]. PCI Security Standards Council. (2024). *Payment Card Industry Data Security Standard: Re- quirements and Testing Procedures* (Version 4.0.1). Retrieved from https://www.pcisecuritystandards. org/document_library

[4]. PCI Security Standards Council. (n.d.). *PCI Security Standards Council releases PCI DSS tokenization guidelines*. Retrieved from https://www.pcisecuritystandards.org/about_us/ press_releases/pci-security-standards-council-releases-pci-dss-tokenization-guidelines/

[5]. EMVCo. (n.d.). *EMV payment tokenisation*. Retrieved March 14, 2026, from https://www. emvco.com/emv-technologies/payment-tokenisation/

[6]. EMVCo. (2023). *EMV payment tokenisation: A guide to use cases* (Version 2.2.1). Retrieved
[7]. from https://www.emvco.com/wp-content/uploads/2023/03/EMVCo-Payment-Tokenisation-A-Guide-To-Us 2.1.pdf

[8]. EMVCo. (2022). *EMVCo white paper on payment account reference* (Version 2.1.1). Retrieved from https://www.emvco.com/emv-technologies/payment-tokenisation/

[9]. EMVCo. (2022). *EMV payment tokenisation specification, technical framework FAQ v2.3.1*. Retrieved from https://www.emvco.com/emv-technologies/payment-tokenisation/

[10]. Lopez-Rojas, E. A., Elmir, A., & Axelsson, S. (2016). *PaySim: A financial mobile money simulator for fraud detection*. In *Proceedings of the 28th European Modeling and Simulation Symposium (EMSS 2016)*.

[11]. Lopez-Rojas, E. A. (n.d.). *PaySim: Financial simulator of mobile money service* [Software repository]. GitHub. Retrieved from https://github.com/EdgarLopezPhD/PaySim

[12]. Dal Pozzolo, A., Caelen, O., Johnson, R. A., & Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced

classification. In *2015 IEEE Symposium Series on Computational Intelligence (SSCI)* (pp. 159–166). IEEE. https://doi.org/10.1109/SSCI.2015.33

[13]. Dal Pozzolo, A., Caelen, O., Le Borgne, Y.-A., Waterschoot, S., & Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Ap- plications, 41* (10), 4915–4928. https://doi.org/10.1016/j.eswa.2014.02.026

[14]. Machine Learning Group, Université Libre de Bruxelles, & Worldline. (2016). *Credit card fraud detection* [Data set]. Retrieved from https://berd-platform.de/records/qcqqe-g6q16

[15]. Seeja, K. R., & Zareapoor, M. (2014). FraudMiner: A novel credit card fraud detection model based on frequent itemset mining. *The Scientific World Journal, 2014*, 252797. https://doi. org/10.1155/2014/252797

[16]. Chung, J., & Lee, K. (2023). Credit card fraud detection: An improved strategy for high recall using KNN, LDA, and linear regression. *Sensors, 23* (18), 7788. https://doi.org/10.3390/ s23187788

[17]. Btoush, E. A. L. M., Zhou, X., Gururajan, R., Chan, K. C., Genrich, R., & Sankaran, P. (2023). A systematic review of literature on credit card cyber fraud detection using machine and deep learning. *PeerJ Computer Science, 9*, e1278. https://doi.org/10.7717/peerj-cs.1278

[18]. Flamini, A., Sciarretta, G., Scuro, M., Sharif, A., Tomasi, A., & Ranise, S. (2024). On crypto- graphic mechanisms for the selective disclosure of verifiable credentials. *Journal of Information Security and Applications, 83*, 103789. https://doi.org/10.1016/j.jisa.2024.103789

[19]. Kahn, C. M., McAndrews, J. J., & Roberds, W. (2005). Money is privacy. *International Eco- nomic Review, 46* (2), 377–399. https://doi.org/10.1111/j.1468-2354.2005.00323.x

[20]. Camenisch, J., Drijvers, M., & Lehmann, A. (2016). Anonymous attestation using the strong Diffie-Hellman assumption revisited. In *Trust and Trustworthy Computing* (pp. 1–20). Springer. https://doi.org/10.1007/978-3-319-45572-3_1

[21]. Islam, M. M., & In, H. P. (2024). An auditable, privacy-preserving, transparent unspent trans- action output model for blockchain-based central bank digital currency. *IEEE Open Journal of the Computer Society, 5*, 671–683. https://doi.org/10.1109/OJCS.2024.3486193

[22]. Liu, T., Liu, Y., Zhang, D., Chen, C., & Wang, W. (2023). ASOZ: A decentralized payment system with privacy preserving and auditing on public blockchain. *Cryptology ePrint Archive*, Paper 2023/1816. Retrieved from https://eprint.iacr.org/2023/1816

[23]. Davidow, D. M., Manevich, Y., & Toch, E. (2023). Privacy-preserving payment system with verifiable local differential privacy. *Cryptology ePrint Archive*, Paper 2023/126. Retrieved from https://eprint.iacr.org/2023/126