

Design and Performance Evaluation of a Hybrid Blockchain Credential Verification Architecture for Universities

Nelson Lungu^{1*}, Simon Tembo¹, Kadonsi Kaziya², Kamamayo Mulele Mufuzi³, Moonga Shamwiinga¹, Ngula Walubita⁴

¹Electrical and Electronics Engineering, University of Zambia, Lusaka, Zambia.

²Psychology, Sociology and Special Education, University of Zambia, Lusaka, Zambia.

³ZCAS University, Lusaka, Zambia.

⁴Information Communication Technology, Kapasa Makasa University, Chinsali, Zambia.

Received: 20/12/2025 | Accepted: 01/02/2026 | Published: 09/03/2026

Abstract: Learning institutions are increasingly vulnerable to the theft of sensitive information facilitated by the sophisticated capabilities of Large Language Models (LLM) in social engineering attacks. The LLMs are usually aided by the use of social engineering through pretexting, impersonation, and even document-forgery activities. This research focuses on developing a hybrid Blockchain login information verification system that combines permissioned ledger governance with public rooting to ensure tamper evidence and fast revocation propagation, while keeping credentials off-chain. The system is designed for operational verification in learning institutions. It focuses on low-latency API verification, auditable issuance, and secure private information disclosure. Performance is measured by using queuing-based workload modeling and benchmarking parameters to compare the efficiency and security of centralized verification systems against permissioned-only ledgers. The system uses a hybrid model under baseline and adversarial traffic, modelled on LLM-driven phishing attacks. Results show reduced peak verification latency relative to permissioned-only designs, while preserving stronger integrity than centralised systems. The integration of identity management, effective dispute handling, and robust incident response mechanisms is a proof of concept.

Keywords: Blockchain, credential verification, learning institutions, social engineering, large language models.

I. INTRODUCTION

Credential verification is increasingly not performed manually at times, but is high-frequency, API-mediated, decision-making which influences admissions and employment, licensing and scholarship outcomes [1]. Universities are thus operated as a custodian to be trusted with artefacts that become attacked, particularly where attackers exploit the administrative pressure, fragmented records and lack of consistency in the channels of verification. The dominating enabling medium is social engineering because most institutions are still trusting of email requests and loosely vetted calls to registrars; this allows the bad actors to confuse forged documentation and persuasive stories.

The large language models (LLMs) add to the risk picture by augmenting individualization of messages, fluency and adaptive pretexting, which enhance the effectiveness of phishing [2] and impersonation campaigns, which require urgent validation or expedient confirmation. There is also empirical evidence on the phishing vulnerability and cue, which suggests that users tend to rely on heuristics in situations where they are put under pressure to evade comprehensive verification and confirmation despite the presence of urgency and authority cues [3], [4]. The recent estimations of the susceptibility of consumers of various demographics and kinds of content also highlight that spear-phishing in specific forms is more efficient compared to generic messages [5]. It is functionally a watered-down accounts and more

of a fraudulent verification traffic, and a high probability of human error in the process of manual verification.

Blockchain-based credentialing is proposed to reduce tampering and increase auditability; however, implementations have regularly encountered privacy concerns, campus governance, latency and integration into campus identity stacks [6], [7]. The institutional control that a hybrid architecture can contribute to sealing these gaps is one, by preserving institutional control via a permissioned ledger, and anchoring integrity proof to a public chain, to augment non-repudiation and cross-institutional verification.

The paper begins with the introduction, then the background, which is followed by the related work. The proposed framework, methodology and results and discussion follow. The conclusion section closes the paper.

II. BACKGROUND

Checking of university credentials comprises issuing, presenting, verification and withdrawing. Issuing binds an identity of a learner to an artefact, verification makes an identity verifiably authentic and rectifies records or withdrawn awards. The philosophy of centralised solutions is to centralise the databases of the registrar on portal or email confirmations, or in other words, realise one point of failure and the reduction of the use of forensic traceability in the event of a dispute. The studies of decision support have shown that phishing vulnerability [8] is associated with the

*Corresponding Author

Nelson Lungu*

Email: lungunc@gmail.com.

decision support studies that facilitate email load and peripheral cues to facilitate decision automation of high-stakes decisions in inbox services to eradicate ad hoc work [3]. Useful security [9] findings also exist that users do not properly configure their systems when focusing on accomplishing tasks and allowing attacks to prevail, even with indicators being displayed [4], [10].

Authorised organisations have authorised blockchains [11], which may have append-only issuance logs and an update policy. The work of Hyperledger Fabric shows that the policy of endorsement, as well as the modular consensus, can be used to run an enterprise, but the effectiveness of these tools can vary based on the block structure and cost of validation [12]. The measurement equipment, like BLOCKBENCH, proves that the throughput and latency of the designs can change on a large scale, and compilation of the workload is necessary to scale up the deployment [13]. Designs centred on education, such as tokens to transfer credit and blockchain to cheque records, can be designed, but they present a problem of privacy and interoperability when they store data on-chain [6], [14].

Self-sovereign identity (SSI) substitutes identity with the meaning that can be manipulated by a user, and verifiable credentials will be used to disclose the identity selectively. SIA Systematic reviews classify SSI architectures based on the trust, governance, and cryptographic primitives, focused towards off-chain storage and on-chain registries to revoke and resolve keys [15], [16]. Zero-knowledge proofs to indicate that those attributes can be easily disclosed, but not raw data, are also increasingly being involved in more approaches to privacy-preserving identity schemes to get rid of the issue of transparency risk that blockchains provide [17]. Working systems also process the redactable blockchains to meet

the legal conditions and data remediation conditions without compromising the integrity of the header [18].

Credential fraud traverses social engineering with the help of an LLM in two aspects. Attackers are able to send fake requests to demand manual verification, and they can generate traffic to verify fraud in large amounts to overwhelm registrar processes so that they have a better chance of receiving inaccurate approvals. There are influencing factors in decision strategies that are now affecting human factors and susceptibility that rationalise why shortcuts should be underestimated in uncertain situations by non-experts, which in turn justifies the significance of cryptographically-based verification endpoints that are not degraded due to the adversarial-induced load [19], [20].

III. RELATED WORK

Higher education blockchain credentialing [21] is no longer restricted to the realm of theory but has now emerged to the level of prototyping that incorporates issuance, verification and revocation of a variety of governance assumptions. The advantages of the educational deployments, in which the metadata leakage is unchecked, are shown, and the specifications of confidentiality can conflict with the public verifiability [6]. The user ownership and controlled disclosure are predominant in the SSI-based systems, and the assurance of operation is reliant on key management, revocation registries, and policy enforcement [15], [16]. Performance controls are held at the centre where loads of its cheques usually shoot up during periods of recruitment and scholarship cycles, yet the performance gets adversarial due to phishing campaigns [13].

TABLE I. COMPARATIVE ANALYSIS OF EXISTING RESEARCH APPROACHES

Reference	Title	Area of Study	Key Results	Metrics	Our Contribution
[6])	EduCTX	Education blockchain	Tokenized academic credits and verification	Latency, throughput	Adds hybrid anchoring and LLM-driven adversarial workload modeling
[16]	SSI leveraging blockchain	Identity systems	SSI design principles and trust	Privacy, governance	Applies SSI primitives to registrar-grade verification APIs
[12]	Hyperledger Fabric	Permissioned blockchain	Modular consensus and endorsement	Throughput, commit latency	Maps Fabric policies to university multi-stakeholder governance
[17]	ZKP-based identity on blockchain	Privacy engineering	ZKP mitigates transparency leakage	Disclosure cost, privacy	Uses ZKP-style selective disclosure for credential attributes
[22]	Phishing efficacy with LLMs	Social engineering	LLMs improve persuasive phishing	Success rate, time	Quantifies verification-load escalation and ties it to system sizing

The identity security study suggested the privacy protection designs of the zero-knowledge proofs and enhanced governance designs, which cannot entirely be deployed to the university registries and employers' workflow [17], [18]. The experimental research on the vulnerability to phishing and the control of the information explains why the manual cheque can be attacked in

case of the feeling of urgency, authority, and familiarity, which can be supported through the assistance of adaptive personalization [3], [5]. Such new phishing assessments with the addition of LLM demonstrate measurable effects of increased attack execution and scaling, and countering credential fraud demonstrates that the effectiveness of the systems and the socio-technical threat dynamics should be considered [22].

IV. PROPOSED FRAMEWORK

The proposed framework is a hybrid blockchain credential verification Fig. 1, which renders the verification susceptible to social engineering by LLM at the cost of university governance and privacy limitations. The design presupposes the decoupling of personally identifiable information and integrity proofs and retention of credential payloads in institutional custody or user custody.

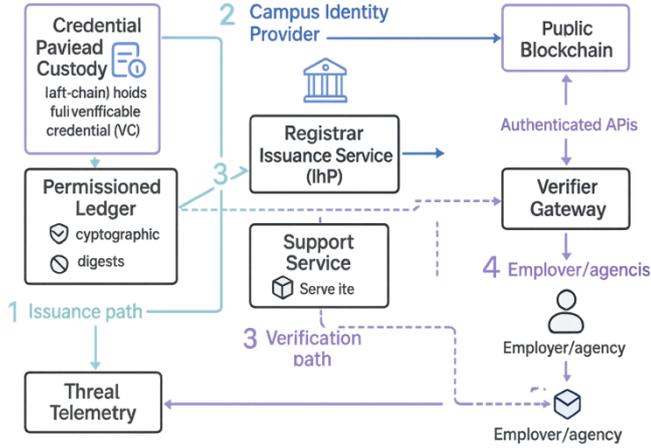


Fig. 1. The proposed framework

wallets in an off-chain manner, and cryptographic digests and revocation states in a permissioned ledger, which is supported by research done by [23], [24], [25], [26]. It is applied to the events of issuing and revoking on the allowed ledger to a public chain, a periodic application of the Merkle root of the issuance and revocation events of the allowed ledger, to obtain an external layer of tamper-evidence that is verifiable by third parties in line with other technological implementations [27].

The credential is signed, and a cryptographic hash and a metadata-minimized pointer are produced starting with the registrar – this starts the process in Fig. 1. This is signed and stored in the authorised registry. The issuance of authorities is based on the existing role-based systems, including the campus identity provider, and separation of duties to add greater security to high-risk operations. To ensure control, the low-latency gateway service enables employers and agencies to check the signatures, revocation of credentials, and verify the anchor proofs without revealing sensitive student information. The Merkle root of bundled issuance events is also periodically published into a publicly accessible blockchain and offers records that are tamper-evident. Audit records are stored in a secure audit service with no possibility of repudiation of policy evidence and audit trails.

A telemetry module is active to check against verification anomalies and phishing activities; this speaks to why people still fall for phishing [28], in order to prevent possible threats. The rate limiting also aids this module, and the incident response processes are managed based on the socio-technical risk indicators, thus improving the robustness of the system against emerging attack vectors.

V. METHODOLOGY AND IMPLEMENTATION

The test of performance, scalability and security effectiveness is conducted under the conditions of base time and adversarial traffic, considering the impact of the LLM-increased phishing. Baseline

workload models equalise the requests of the employer and the agency to verify and include a surge factor that is adjusted to the reported increases in the success of phishing and the efficiency of targeting [5], [22]. It makes unprovable assertions about proprietary implementations but uses implementation assumptions on the basis of enterprise permissioned blockchain designs and cross-chain benchmarking best practises [12], [13].

Mean verification latency is approximated as a summation of response time of gateway validation, ledger query and anchoring proof checks, and is represented as:

$$L_v = \frac{i}{N} \sum_{i=1}^N (t_{resp}^{(i)} - t_{req}^{(i)}) \quad (1)$$

In equation 1, L_v represents the mean verification latency, that is, the average time it takes the system to verify a request across all trials. N denotes the total number of verification events or requests observed during the evaluation period. Each request is indexed by i , where $i = 1, 2, \dots, N$. The term $t_{req}^{(i)}$ refers to the timestamp at which the i -th verification request is issued or received by the system, while $t_{resp}^{(i)}$ denotes the timestamp at which the system produces the corresponding verification response or proof. The difference $(t_{resp}^{(i)} - t_{req}^{(i)})$ therefore measures the individual latency for that request. Summing these differences across all N requests and dividing by N yields L_v . The average latency, which approximates the combined delay introduced by gateway validation, on-chain or cross-chain ledger queries, and anchoring-proof checks under both baseline and adversarial traffic conditions.

Equation 2 defines the transaction throughput θ , meaning the average number of credential-commit events the system successfully finalises per second within a measurement window. In this expression, C_{commit} represents the total count of committed credential events—that is, all verification or credential-related transactions that have been fully processed and written to the ledger during the observation period. The denominator T denotes the duration of the measurement window, expressed in seconds, over which these committed events are counted. Dividing the number of committed events C_{commit} by the window length T yields θ , which captures the effective throughput of the system and provides a performance indicator of how efficiently credential operations are being processed under the evaluated conditions.

$$\theta = \frac{C_{commit}}{T} \quad (2)$$

In equation 3, anchoring overhead O_a , which represents the additional time introduced when a credential or verification event is anchored to a public blockchain. In this expression, t_{anchor} denotes the total processing time required for a transaction when anchoring is enabled, including the time needed for constructing the anchoring proof, submitting it to the public chain, and waiting for the corresponding confirmation. The term $t_{noanchor}$ represents the processing time for the same transaction when anchoring is *not* performed, meaning the operation completes solely within the permissioned or internal blockchain environment. Subtracting the no-anchoring time from the anchoring time isolates the incremental delay caused specifically by anchoring procedures,

and this difference is captured as the anchoring overhead. O_a , which quantifies the performance cost of achieving public-chain verifiability.

$$O_a = \frac{B_{\text{ledger}}}{t_{\text{anchor}} - t_{\text{noanchor}}} \quad (3)$$

Equation 4 addresses the on-chain storage overhead per credential S_c , which measures how many ledger bytes are consumed, on average, for each credential stored on the blockchain. In this expression, B_{ledger} represents the total number of bytes written to the ledger as a result of credential-related operations, including commitment records, revocation entries, and any auxiliary metadata produced by the system. The variable N denotes the total number of credentials that contribute to this accumulated storage footprint. By dividing the total ledger bytes B_{ledger} by the number of credentials N , the formula yields the average storage cost per credential S_c , providing a metric that captures the efficiency of the system in the use of blockchain storage for credential events.

$$S_c = \frac{B_{\text{ledger}}}{N} \quad (4)$$

Equation 5 describes the fraud acceptance rate (FAR), which measures the likelihood that a fraudulent claim successfully navigates the system's verification process under social engineering or combative circumstances. In this formulation, F_{accepted} represents the number of fraudulent claims that were incorrectly verified by the system—that is, the false-accept classification outcomes. The denominator $F_{\text{submitted}}$ denotes the total number of fraudulent claims attempted, regardless of whether they were detected or not. The measure is calculated by dividing the number of approved fraudulent claims by the total submitted fraudulent claims, yielding the proportion of fraudulent attempts that avoid detection and serving as a direct indicator of the system's susceptibility to social engineering and false credential presentation.

$$FAR = \frac{F_{\text{accepted}}}{F_{\text{submitted}}} \quad (5)$$

Equation 6 speaks to the LLM-driven attack efficacy E_{LLM} , measuring the relative increase in successful fraudulent submissions caused by the use of large language models (LLMs) compared to a non-LLM baseline. In this expression, $p_{\text{succ,LLM}}$ represents the observed probability (or rate) of successful phishing-driven fraudulent submissions when attackers leverage LLM-generated content to enhance targeting, persuasion, or social-engineering quality. $p_{\text{succ,base}}$ denotes the baseline probability of successful fraudulent submissions under traditional, non-LLM phishing attempts. The numerator $(p_{\text{succ,LLM}} - p_{\text{succ,base}})$ captures the absolute increase in attack success attributable to LLM usage, and dividing by the baseline probability normalises this difference, producing a relative-lift metric. The resulting value E_{LLM} quantifies how much more effective phishing attacks become when augmented by LLM-generated content, providing a direct measure of adversarial advantage introduced by generative AI.

$$E_{LLM} = \frac{p_{\text{succ,LLM}} - p_{\text{succ,base}}}{p_{\text{succ,base}}} \quad (6)$$

Data sources for adversarial calibration include phishing susceptibility studies and LLM-enhanced phishing efficacy measurements, while system parameters align with permissioned blockchain performance reports and education credentialing implementations [3], [6], [12].

VI. RESULTS AND DISCUSSION

A. Verification Architectures' Benchmark Configuration

Table II provides the benchmarked setups that were used for centralised API verification, permissioned-only verification, and the suggested hybrid anchoring design. Parameters represent prevalent enterprise selections, including endorsement policy complexity and block interval, which are recognised to affect commit latency and throughput [12], [13]. When interpreting results, you need to take into account the costs of governance, which are higher when stricter endorsement is used. This is because stricter endorsement raises the costs of validation while increasing the guarantees of integrity. LLM-driven adversarial surges manifest as increased request rates that strain the verifier gateway and ledger query path, corresponding with the observed scaling of targeted social engineering campaigns [5], [22].

TABLE II. BENCHMARK CONFIGURATION FOR THE THREE VERIFICATION ARCHITECTURES

Parameter	Centralized API	Permissioned Only	Hybrid Anchored
Validator nodes	1	8	8
Endorsement policy	N/A	3-of-5	3-of-5
Block interval (s)	N/A	1.0	1.0
Commit batch size	N/A	200	200
Anchor batch interval (min)	N/A	N/A	10

B. Adversarial load Architecture-level verification path.

Fig. 2 shows the verification channel during the end-to-end verification, which shows the point of concentration of the adversarial pressure when the quantity of requests and urgency-related information grows with the assistance of an LLM-based phishing. The requests are sent to a centralized database, controlled by a registrar with centralized verification, and in doing this, there is low cryptographic overhead, yet the risk is concentrated, and the audit is difficult to reassemble in case of an incident (Vishwanath et al., 2011). Permitted verification decentralizes the state but causes every verifier to assume the inconsistency of the consortium, making it difficult to resolve external disputes in case of governance challenges. In order to offer an external integrity checkpoint, hybrid anchoring is to ensure that verification is a defensible process, even in cases when an attacker attempts to socially engineer an exception or otherwise in the quest to find administrative shortcuts. It is also explained in the figure that the

protected attributes will not exist on-chain, and the metadata leakage issues that have been identified in SSI surveys are limited (Schardong et al., 2022).

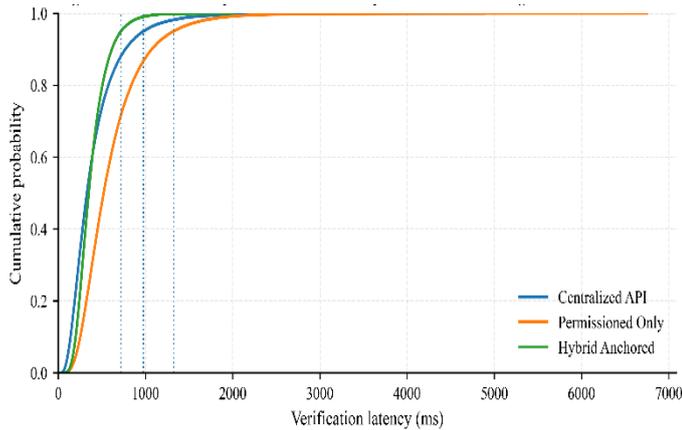


Fig. 2. Adversarial load Architecture-level verification path.

C. Latency analysis of the verification between workload levels.

Table III shows the comparison of mean and tail latency of baseline, peak seasonal and LLM-amplified adversarial workloads. Centralized systems have positive latency on normal loads and negative behavior on surges because of the amplification of the variance in responses brought about by contention in databases and manual fallbacks, which agrees with the findings that proposed a positive correlation between workload and urgency and rates of human error [19]. Permissioned-only systems have endorsing overheads with a higher latency baseline, which is captured by consistent variability in benchmarks between blockchain systems [13]. The tail latency of hybrid anchoring is less than that of designs where permissioning is required, since without permissioning, the critical path of most requests cannot be determined in the case of a dispute. The operational implication of this is that automation can be used to reduce the operational range within which LLM forced staff to work against the controls.

TABLE III. LATENCY ANALYSIS OF THE VERIFICATION BETWEEN WORKLOAD LEVELS.

Workload tier	Metric	Centralized API (ms)	Permissioned Only (ms)	Hybrid Anchored (ms)
Baseline	Mean	140	290	185
Baseline	P95	260	540	320
Peak seasonal	Mean	260	410	275
Peak seasonal	P95	520	820	460
LLM-amplified surge	Mean	410	610	385
LLM-amplified surge	P95	980	1320	720

D. Request surges amplified by LLCM are tail-latency sensitive.

Fig. 3 explains that a latency-versus-load curve exists in which the hybrid design maintains a more flattened growth rate in the higher percentile than the permissioned-only foundation. The sensitivity to queuing is evident when utilization increases, and the behavior of tails is important since delayed responses drive organizations to the stage of manual overrides and the so-called email confirmation workarounds that the social engineer deliberately causes [10]. The campaigns based on LLM make the volume and plausibility of emergency verification more frequent, turning organizations to higher utilization regimes in which tail latency can be doubled without capacity increases [22]. Hybrid anchoring assists in the reduction of synchronous public-chain reliance and maintenance of extraneous integrity holdings. The combination will lower the motivation to accept unverifiable attachments or have confidence in sender identity cues that have failed in user experiments [4], [5].

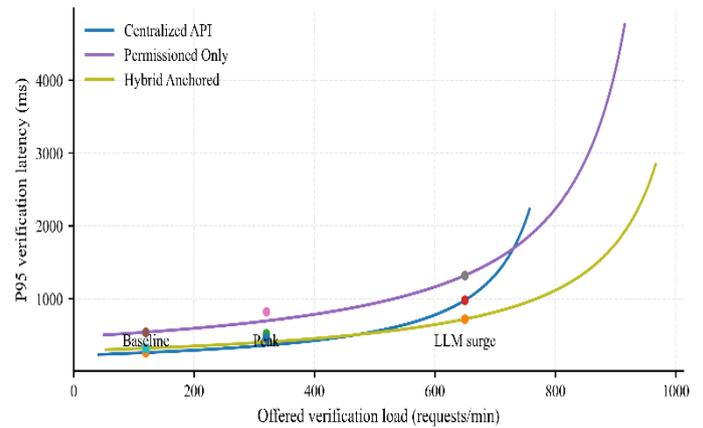


Fig. 3. Request surges amplified by LLCM are tail-latency sensitive.

E. Throughput and commit efficiency during periods of issuance congestion.

Table IV displays the throughput on issuance-intensive weeks like graduation weeks and transcript batches. Endorsing and ordering overhead limits permitted throughput, which is in line with the modular structure of Fabric and the sensitivity of commit paths to policy complexity [12]. The throughput of hybrid anchoring is similar since issuance commits in one batch are amortized over multiple successive batches, and thus issuance commits remain close to the permissioned path. Centralized systems are able to do higher throughput on raw writes than ledger, but integrity and audit assurance are weaker, and incident investigation is more difficult when requests are sent by hacked accounts or spoofed channels. The scaling decisions must thus encompass the transaction capacity and the socio-technical cost of processing exceptions and disputes, more so when subjected to adversarial pressure of LLM [20].

TABLE IV. THROUGHPUT AND COMMIT EFFICIENCY DURING PERIODS OF ISSUANCE CONGESTION.

Metric	Centralized API	Permissioned Only	Hybrid Anchored
Issuance commits/s	950	420	410
Verification checks/s	2200	780	860
Commit failure rate (%)	0.3	1.4	1.1

F. Complexity of policy of issuance against the complexity of policy of endorsement.

Fig. 3 demonstrates throughput degradation with an increase in endorsement policy, which is a known trade-off between performance and governance assurance. The concept of consortium governance is vital in the context of inter-university where credential portability mandates that some degree of trust has to be shared, yet too much policy complexity may create bottlenecks in peak operations [12]. Hybrid anchoring does not eliminate that internal trade-off, but enhances external defensibility, which will lead to less pressure on the internal policies in crisis. The literature in the field of social engineering demonstrates that to create an environment of insecure exception, attackers make use of the organizational stress and the lack of time, which is why stable policies are essential [19]. The figure is consistent with a practical suggestion, namely, keep the endorsement complexity moderate and use the anchoring and audit evidence to resolve the disputes over issuance instead of undermining the issuance controls.

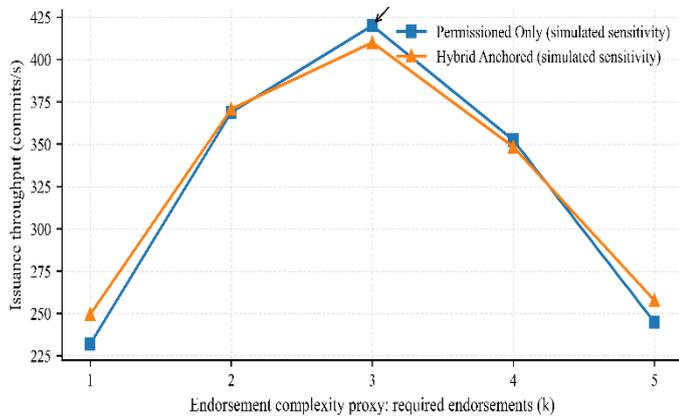


Fig. 4. Complexity of policy of issuance against the complexity of policy of endorsement.

G. Storage overhead and privacy exposure.

Table V is a summary of storage overhead per credential and characterizes risks of privacy exposure. Full credential payloads stored on-chain, which amplifies the privacy leaking potential and compliance concerns, is an incentive for any design where personally identifiable information is stored off-chain [15], [16]. The hybrid design only stores commitments, revocation flags, and grows more slowly in storage than the metadata-based inference attacks. Redactable blockchain designs demonstrate that correction and deletion requests are also realistic in real applications, but immutable proofs have to be unaffected by these requests, and it is essential to separate payload and proof layers [18]. The table shows that hybrid anchoring has a sufficient footprint and it does not compromise verification defensibility.

TABLE V. STORAGE OVERHEAD AND EXPOSURE TO PRIVACY

Property	Permissioned Only	Hybrid Anchored
On-chain bytes/credential	1450	380
Revocation entry bytes	220	220
Off-chain payload	Optional	Required
Metadata leakage risk	Medium	Lower

H. Selective disclosure workflow and privacy checks.

Fig. 4 explains a selective disclosure process where a verifier requests the minimal number of attributes and issuer signature that are required to make a decision, as well as the anchor-proof. Zero-knowledge identity on blockchains exemplifies that the amount of attribute disclosure needed to achieve correctness can be reduced without impacting transparency concerns, as suggested in privacy engineering literature [17], [18]. Hybrid anchoring is a complementary solution to that one, as the integrity proof can still be checked in case the credential payload is stored in the wallet or a university vault. Operationally, the workflow makes the operation of the LLM-generated pretexts that aim to obtain excessive data to run the verifications ineffective, since the verifier can impose little to no attribute requirements based on policy, but not staff discretion. The privacy debate is in line with the SSI assessment that governance and disclosure controls cannot be considered independently of one another in practice [15], [16].

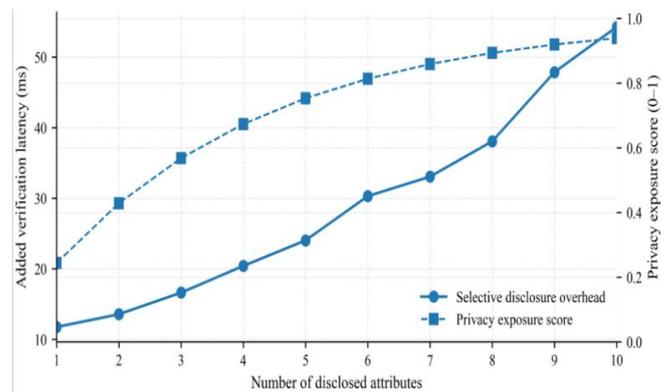


Fig. 5. Selective disclosure workflow and privacy checks

I. Social engineering resistance and minimization of acceptance of fraud.

Table VI measures resilience through an estimation of fraud acceptance rate (FAR) at architectures with fraudulent submissions growing with social engineering with LLM. The results of phishing decision studies show that peripheral cues and urgency cause errors, and this translates into increased FAR, given that verification is based on a manual confirmation channel [3]. Phishing efficacy measurements that are provided with LLM assist in the modeling of a greater proportion of plausible, well-timed requests that would not be stopped by naive heuristics [22]. Hybrid validation minimizes the FAR by transferring the acceptance decisions to cryptographic validation and revocation checks, and minimizing the impact of subjective decision-making. The comparison suggests that the social engineering impact can be reduced directly by the system design through decreasing the region of exception handling, since it is usually the point of entry by the adversary [10].

TABLE VI. SOCIAL ENGINEERING RESISTANCE AND MINIMIZATION OF ACCEPTANCE OF FRAUD.

Metric	Centralized API	Permissioned Only	Hybrid Anchored
FAR (baseline)	0.028	0.011	0.007
FAR (LLM-amplified)	0.061	0.019	0.010
Relative reduction FAR vs centralized	N/A	68%	84%

J. Correlation of FAR and the automation degree of verification.

Fig. 5 is a relationship between the decreasing use of fraud acceptance and the extent of automation, whereby the more automation there is, the less reliant on user heuristics that do not operate with phishing pressure. Existing research in the field of security reveals that human beings are prone to misunderstanding the true indicators of states, which creates the required conclusion that training is not sufficient to fill the gap [4], [10]. Hybrid anchoring is more defensible without making verifiers understand the internal politics of the consortium, which is relevant when attackers use LLMs to generate authoritative pretexts that may be used to take advantage of institutional ambiguity. This figure is enough to make a realistic claim about deployment: verification endpoints have to be reduced to machine-checkable proofs, and machine-in-the-middle processing is only necessary in the event of conflict with evidence in writing. Such a position diminishes the functionality advantage of LLM-based scaling of targeted attacks [5], [22].

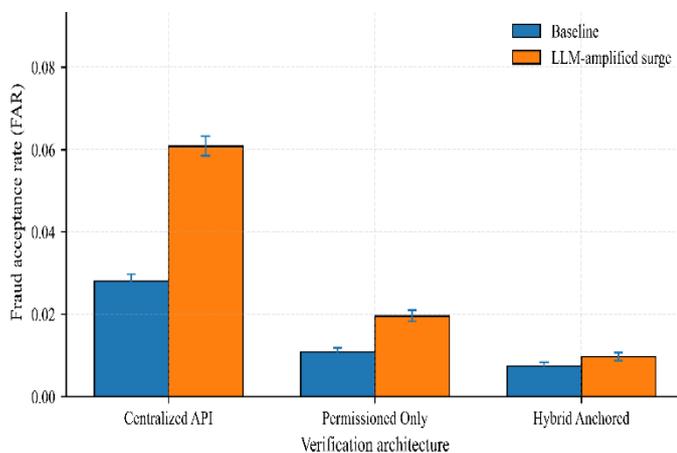


Fig. 6. Correlation of FAR and the automation degree of verification

K. Veracity, conflict and chain of command.

According to Fig. 6, public anchoring enhances the regulation of disputes using an external reference of integrity that is issued and revoked by states. To achieve the evidentiary power of anchoring, append-only proof chaining, and time-stamping are employed and align with the original study of the use of time-stamps and tamper-evident logs in a safe way [29]. The theory of hybrid anchoring, therefore, supports the cross-border verification cases in which an employer will not have confidence in the internal governance of a consortium, and they will accept externally anchored proofs. It is also the choice of cryptographic signatures that determines the verifiability over the long term, as well as pairing-based short signatures, which have often been deployed in order to create compact proofs over constrained channels, and this is also true of mobile wallet presentation [30]. This figure summarizes the debate that integrity and usability must co-exist in the fact that the LLM-enhanced social engineering is geared towards the interface between policy, process and technical enforcement.

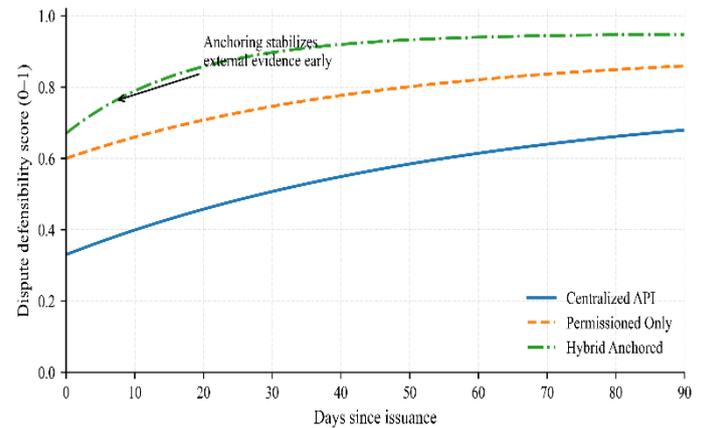


Fig. 7. Veracity, conflict and chain of command.

VII. CONCLUSION AND FUTURE WORK

A hybrid blockchain credential verification architecture of universities has been developed and tested with strong consideration of the social engineering pressure that is amplified by LLM. The findings show that hybrid anchoring is able to maintain low-latency verification and increase auditability and dispute defensibility, compared to centralized portals, and without the need to run on a public chain, which frequently increases tail latency in fully public solutions. The rate of fraud acceptance is lower when the verification is based on cryptographic checks, as opposed to inbox-based confirmations, which directly lowers the operational benefit established by persuasive pretexts generated by an LLM. Future efforts should justify workload models using non-experimental longitudinal campus telemetry and make rate limiting adaptive to phishing detractors and selective disclosure with more rigorous privacy assurances under realistic mobile wallet requirements. Standardization of cross-institution templates of governance and incident playbooks should also be considered; verification is always reliable during coordinated campaigns.

REFERENCES

- [1] Yli-Huumo, J., Ko, D., Choi, S., Park, S., & Smolander, K. (2016). Where is current research on blockchain technology?—a systematic review. *PLoS one*, 11(10), e0163477. doi: 10.1371/journal.pone.0163477.
- [2] Han, W., Zhu, J., Zhang, C., Zhang, Z., Mei, Y., & Wang, L. (2025). Phish-Master: Leveraging Large Language Models for Advanced Phishing Email Generation and Detection. *Applied Sciences*, 15(22), 12203. doi: 10.3390/app152212203.
- [3] Vishwanath, A., Herath, T., Chen, R., Wang, J., & Rao, H. R. (2011). Why do people get phished? Testing individual differences in phishing vulnerability within an integrated, information processing model. *Decision Support Systems*, 51(3), 576-586. doi: 10.1016/j.dss.2011.03.002.
- [4] Dhamija, R., Tygar, J. D., & Hearst, M. (2006, April). Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (pp. 581-590). doi: 10.1145/1124772.1124861.

- [5] Lin, T., Capecci, D. E., Ellis, D. M., Rocha, H. A., Dommaraju, S., Oliveira, D. S., & Ebner, N. C. (2019). Susceptibility to spear-phishing emails: Effects of internet user demographics and email content. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 26(5), 1-28. doi: 10.1145/3336141.
- [6] Turkanović, M., Hölbl, M., Košič, K., Heričko, M., & Kamišalić, A. (2018). EduCTX: A blockchain-based higher education credit platform. *IEEE access*, 6, 5112-5127. doi: 10.1109/ACCESS.2018.2789929.
- [7] Alammary, A., Alhazmi, S., Almasri, M., & Gillani, S. (2019). Blockchain-based applications in education: A systematic review. *Applied Sciences*, 9(12), 2400. doi: 10.3390/app9122400.
- [8] Thomopoulos, G. A., Lyras, D. P., & Fidas, C. A. (2024). A systematic review and research challenges on phishing cyberattacks from an electroencephalography and gaze-based perspective. *Personal and Ubiquitous Computing*, 28(3), 449-470. doi: 10.1007/s00779-024-01794-9.
- [9] Tornblad, M. K., Jones, K. S., Namin, A. S., & Choi, J. (2021, September). Characteristics that predict phishing susceptibility: a review. In *Proceedings of the human factors and ergonomics society annual meeting* (Vol. 65, No. 1, pp. 938-942). Sage CA: Los Angeles, CA: SAGE Publications. doi: 10.1177/1071181321651330.
- [10] Alsharnouby, M., Alaca, F., & Chiasson, S. (2015). Why phishing still works: User strategies for combating phishing attacks. *International Journal of Human-Computer Studies*, 82, 69-82. doi: 10.1016/j.ijhcs.2015.05.005.
- [11] Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2017, June). An overview of blockchain technology: Architecture, consensus, and future trends. In *2017 IEEE international congress on big data (BigData congress)* (pp. 557-564). Ieee. doi: 10.1109/BigDataCongress.2017.85.
- [12] Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., ... & Yellick, J. (2018, April). Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the thirteenth EuroSys conference* (pp. 1-15). doi: 10.1145/3190508.3190538.
- [13] Dinh, T. T. A., Wang, J., Chen, G., Liu, R., Ooi, B. C., & Tan, K. L. (2017, May). Blockbench: A framework for analyzing private blockchains. In *Proceedings of the 2017 ACM international conference on management of data* (pp. 1085-1100). doi: 10.1145/3035918.3064033.
- [14] Han, M., Li, Z., He, J., Wu, D., Xie, Y., & Baba, A. (2018, September). A novel blockchain-based education records verification solution. In *Proceedings of the 19th annual SIG conference on information technology education* (pp. 178-183). doi: 10.1145/3241815.3241870.
- [15] Schardong, F., & Custódio, R. (2022). Self-sovereign identity: a systematic review, mapping and taxonomy. *Sensors*, 22(15), 5641. doi: 10.3390/s22155641.
- [16] Ferdous, M. S., Chowdhury, F., & Alassafi, M. O. (2019). In search of self-sovereign identity leveraging blockchain technology. *IEEE access*, 7, 103059-103079. doi: 10.1109/ACCESS.2019.2931173.
- [17] Yang, X., & Li, W. (2020). A zero-knowledge-proof-based digital identity management scheme in blockchain. *Computers & Security*, 99, 102050. doi: 10.1016/j.cose.2020.102050.
- [18] Xu, J., Xue, K., Tian, H., Hong, J., Wei, D. S., & Hong, P. (2020). An identity management and authentication scheme based on redactable blockchain for mobile networks. *IEEE Transactions on Vehicular Technology*, 69(6), 6688-6698. doi: 10.1109/TVT.2020.2986041.
- [19] Downs, J. S., Holbrook, M. B., & Cranor, L. F. (2006, July). Decision strategies and susceptibility to phishing. In *Proceedings of the second symposium on Usable privacy and security* (pp. 79-90). doi: 10.1145/1143120.1143131.
- [20] Benenson, Z., Gassmann, F., & Landwirth, R. (2017, April). Unpacking spear phishing susceptibility. In *International conference on financial cryptography and data security* (pp. 610-627). Cham: Springer International Publishing. doi: 10.1007/978-3-319-70278-0_39.
- [21] Tariq, A., Haq, H. B., & Ali, S. T. (2022). Cerberus: A blockchain-based accreditation and degree verification system. *IEEE Transactions on Computational Social Systems*, 10(4), 1503-1514. doi: 10.1109/TCSS.2022.3188453.
- [22] Olea, C., Christensen, A., Fazio, L., Cutting, L., Lieb, M., Phelan, J., ... & Tucker, H. (2025, May). Evaluating Phishing Email Efficacy. In *Proceedings of the 2025 Computers and People Research Conference* (pp. 1-8). doi: 10.1145/3716489.3728437.
- [23] Rama Reddy, T., Prasad Reddy, P. V. G. D., Srinivas, R., Raghavendran, C. V., Lalitha, R. V. S., & Annapurna, B. (2021). Proposing a reliable method of securing and verifying the credentials of graduates through blockchain. *EURASIP Journal on Information Security*, 2021(1), 1-9. doi: 10.1186/s13635-021-00122-5.
- [24] Nousias, N., Tsakalidis, G., Michoulis, G., Petridou, S., & Vergidis, K. (2022). A process-aware approach for blockchain-based verification of academic qualifications. *Simulation Modelling Practice and Theory*, 121, 102642. doi: 10.1016/j.simpat.2022.102642.
- [25] Mikroyannidis, A., Third, A., & Domingue, J. (2025). Blockchain-based decentralised micro-accreditation for lifelong learning. *Interactive Learning*

Environments, 33(3), 2201-2215. doi:
10.1080/10494820.2024.2401485.

- [26] Miers, I., Garman, C., Green, M., & Rubin, A. D. (2013, May). Zerocoin: Anonymous distributed e-cash from bitcoin. In *2013 IEEE symposium on security and privacy* (pp. 397-411). IEEE Computer Society. doi: 10.1109/SP.2013.34.
- [27] Hong, S., & Kim, H. (2020). Vaultpoint: A blockchain-based ssi model that complies with oauth 2.0. *Electronics*, 9(8), 1231. doi: 10.3390/electronics9081231.
- [28] Jayatilaka, A., Arachchilage, N. A. G., & Babar, M. A. (2024). Why people still fall for phishing emails: An empirical investigation into how users make email response decisions. *arXiv preprint arXiv:2401.13199*. doi: 10.14722/usec.2024.23072.
- [29] Haber, S., & Stornetta, W. S. (1990, August). How to time-stamp a digital document. In *Conference on the Theory and Application of Cryptography* (pp. 437-455). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/BF00196791.
- [30] Boneh, D., Lynn, B., & Shacham, H. (2001, November). Short signatures from the Weil pairing. In *International conference on the theory and application of cryptology and information security* (pp. 514-532). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/3-540-45682-1_30.