



Beyond Random Suffixes: Security Analysis and a Cryptographically Hardened Reference-Number Protocol for Web-Based Academic Credential Verification, with a Case Implementation at Northwest University, Kano

Faruk Aliyu Fari¹, Yusuf Aliyu Adamu^{2*}, Usman Mahmud³, Sani Bature Sufyan⁴, Fatima Shehu Tofa⁵, Abdullahi Abdulwahab⁶

¹Department of Computer Science, Faculty of Computing, Northwest University, Kano, Nigeria

²Department of Information and Communication Technology, Faculty of Computing, Northwest University, Kano, Nigeria

³Department of Software Engineering, Faculty of Computing, Northwest University, Kano, Nigeria

⁴Department of Computer Science, Faculty of Computing, Northwest University, Kano, Nigeria

⁵Department of Computer Science, Faculty of Science and Computing, Baba-Ahmed University, Kano, Nigeria

⁶Department of Cyber and Information Technology, Faculty of Computing, Bayero University, Kano, Nigeria.

*Corresponding Author

Yusuf Aliyu Adamu

Article History

Received: 04.06.2026
Revised: 13.07.2026
Accepted: 14.08.2026
Published: 10.09.2026

This work is licensed under a [Creative Commons Attribution-NonCommercial 4.0 International License](https://creativecommons.org/licenses/by-nc/4.0/) (CC BY-NC 4.0)

Abstract: Many institution-level credential-verification systems generate a random-looking reference code and treat it as secure by virtue of being unique, without examining whether that code is actually resistant to guessing. This paper re-examines a deployed reference-number scheme built for the Academic Document Verification System at Northwest University (NWU), Kano, and shows that its uniqueness suffix - four hexadecimal characters drawn from a UUID4 - carries only 16 bits of entropy (65,536 possible values). We show analytically and by Monte Carlo simulation that, while the system's retry-on-collision loop guarantees no reference is ever issued twice, the same 16-bit space can be exhaustively enumerated against the public verification endpoint in under 11 minutes at a modest 100 requests per second, since no rate limiting is applied. We then design, specify, and analyse a drop-in replacement: an HMAC-SHA256-derived tag computed over the year, document-type, department, and a monotonic serial number, keyed by a server-side secret. At a 48-bit truncation this raises brute-force exhaustion time to an estimated 89,000 years under the same attack rate, at a measured computational cost of 3.2 microseconds per generation or verification - statistically negligible against the system's sub-second query latency. The two schemes are compared against manual, QR-code-based, blockchain-anchored, and zero-knowledge credential-verification approaches reported in the literature on a five-dimension feature matrix, supplemented by a quantitative verification-time comparison. The system was validated through 34 functional test cases (100% pass rate) and an exploratory, formative user-acceptance evaluation (n=8; overall satisfaction 4.5-4.8/5.0, not powered for statistical inference), in which security confidence was the lowest-rated criterion - a finding consistent with, though not statistically proof of, the entropy weakness this paper identifies and addresses. The contribution of this paper is therefore not the verification system as an artefact, but a transferable method for quantifying and closing the guessability gap in identifier-based credential-verification schemes generally.

Keywords: credential verification; reference-number entropy; HMAC; brute-force resistance; birthday bound; academic document security.

Cite this article:

Fari, F. A., Adamu, Y. A., Mahmud, U., Sufyan, S. B., Tofa, F. S., Abdulwahab, A., (2026). Beyond Random Suffixes: Security Analysis and a Cryptographically Hardened Reference-Number Protocol for Web-Based Academic Credential Verification, with a Case Implementation at Northwest University, Kano. *ISAR Journal of Science and Technology*, 4(9), 33-39.

1.0 Introduction

Institutions adopting web-based academic-credential verification frequently generate a 'unique reference number' for each document and assume that uniqueness alone confers security. Uniqueness and unguessability are not the same property: a reference number can be perfectly unique within a database while still being drawn from a keyspace small enough to be enumerated by an automated script against a public-facing endpoint. This distinction is rarely made explicit in the applied literature on Nigerian university information systems, where reported systems (Adewale et al., 2020; Olayemi et al., 2019) are evaluated for functional correctness and user acceptance but not, typically, for the entropy of the identifier itself or the cost of attacking it.

This paper uses the Academic Document Verification System built for Northwest University (NWU), Kano, as a case vehicle for closing that gap. The system's as-built reference number is structured as NWU-`{year}`-`{doctype}`-`{dept}`-`{suffix}`, where the suffix is the first four hexadecimal characters of a UUID4. We first show, both by closed-form probability and by simulation, that this suffix carries only 16 bits of entropy and is trivially enumerable once an attacker exploits the fact that the surrounding fields (year, document type, department) are publicly documented rather than secret. We then propose, specify, and analyse a keyed HMAC-based replacement, and quantify its cost and benefit relative to the baseline and to QR-code- and blockchain-based alternatives reported elsewhere.

The specific objectives of this paper are to: (i) formally characterise the entropy and collision behaviour of the baseline reference-number scheme; (ii) quantify the practical feasibility of a brute-force enumeration attack against it; (iii) design a cryptographically keyed alternative that closes this gap without materially increasing computational cost or infrastructure requirements; and (iv) empirically validate both the functional system and the proposed hardening through structured testing, simulation, and user-acceptance evaluation.

2.0 Related Work

Heeks (2006) frames institutional digitisation, including credential verification, within the wider e-government adoption literature, noting that trust in a digital system is as much a function of its demonstrable security properties as of its usability. Okonkwo and Okafor (2020) and NITDA (2021) document the scale of academic credential fraud in Nigeria and argue for centralised digital verification infrastructure, but neither examines the identifier design of the systems they survey. Adewale et al. (2020) demonstrate a QR-code-based certificate verification system; the QR code itself typically encodes a URL or serial number whose guessability is bounded by the length of the encoded serial,

but the approach ties verification to a printed, scannable artefact that can degrade or be excluded from a photocopy. Boonkrong (2024) approaches the same forgery problem from the document side, proposing an automated academic-document forgery-detection system that inspects the presented artefact itself rather than checking it against an issued identifier, a complementary rather than competing strategy to reference-number-based verification.

A substantial and more recent strand of the literature pursues blockchain- and zero-knowledge-based credentialing. Alsobhi et al. (2023), in a systematic review of blockchain-based micro-credentialing in higher education, note that while such systems offer strong tamper-evidence, institutional adoption is frequently constrained by ledger infrastructure, node-maintenance, and interoperability costs - the same trade-off Tapscott and Tapscott (2016) identify at a general level. Xu (2024) reports a working blockchain-based academic credential verification system built specifically for a university-adoption context, while Ramić et al. (2024) and Berrios Moya et al. (2025) extend this line of work toward zero-knowledge and selective-disclosure credential schemes that let a holder prove specific claims (for example, degree class) without revealing the full underlying record. These approaches address a genuine and different problem - selective disclosure and tamper-evident anchoring - from the one this paper addresses, which is the narrower question of whether a lightweight, non-blockchain reference identifier can itself resist guessing; the two lines of work are complementary rather than competing, and a future NWU deployment could in principle layer selective-disclosure guarantees on top of the keyed reference scheme proposed here.

Within applied cryptography, the security of a keyed message authentication code such as HMAC is formally established by Krawczyk et al. (1997) and treated at length by Menezes et al. (1996), who show that an adversary without the key gains no computational advantage over exhaustive search of the tag space, provided the underlying hash function behaves as assumed. The specific failure mode this paper analyses - a reference or token whose 'uniqueness' masks low guessing entropy - is a recognised and formally catalogued weakness class (CWE-331: Insufficient Entropy; MITRE, 2024) and is explicitly flagged for HMAC-based tokens by current IETF guidance on JSON Web Token best practice (Sheffer et al., 2020), which warns that a token's security is bounded by the entropy of its secret and tag space regardless of how cryptographically respectable the surrounding construction looks. Danish et al. (2026) provide the most direct empirical grounding for the risk modelled analytically in this paper: in a large-scale measurement study of SMS-delivered access tokens across more than 300,000 real URLs, they found 125 production services whose tokens were enumerable in

practice specifically because of low token entropy, despite each token being unique per user. This paper connects that formally catalogued and empirically observed weakness class to the applied, institution-level credential-verification literature represented by Adewale et al. (2020) and Okonkwo and Okafor (2020): where the latter treats 'unique reference number' as a sufficient security property, we show, analytically and by simulation, that it is necessary but not sufficient, and provide the keyed-MAC construction that closes the difference at negligible computational cost.

3.0 Materials and Methods

3.1 System Overview

The verification system follows a three-tier architecture: Bootstrap 5/Jinja2 server-rendered templates (presentation), a Flask application handling authentication, routing, and reference-number logic (application), and SQLAlchemy over SQLite for persistence (data), with all queries parameterised through the ORM (Grinberg, 2018; Hipp, 2023). The functional surface - administrator CRUD over document records and a public single-field verification portal - follows the Waterfall model (Royce, 1970) through requirements, design, implementation, and test phases (Pressman & Maxim, 2015). This functional layer is not this paper's contribution and is described only to the extent needed to situate the analysis that follows.

3.2 Baseline Reference-Number Scheme

Table 1 decomposes the as-built reference number. Four of its five components - institution code, year, document-type code, and department code - are derived from public or semi-public information about the document and are not intended to be secret. Only the fifth component, a four-character hexadecimal suffix, is intended to make the reference number 'unpredictable'.

Component	Position	Example	Derivation Rule
Institution code	1st	NWU	Fixed institutional prefix
Issuance year	2nd	2026	Extracted from date_issued
Document-type code	3rd	TRNS/STMT/PUTME/ADMSN	Lookup dictionary on document_type
Department code	4th	CSC	Characters [4:7] of matric number
Uniqueness suffix	5th	A1B2	First 4 hex chars of UUID4; retried on collision

Algorithm 1: Baseline reference-number generation (as-built)

```

Input: matric_number, document_type, date_issued
1 year <- date_issued.split('-')[0]
2 dept <- upper(matric_number[4:7])
3 code <- DOC_TYPE_MAP[document_type]
4 repeat
5   suffix <- upper(hex(uuid4())[0:4]) // 16 bits
6   ref <- "NWU-" + year + "-" + code + "-" + dept + "-" + suffix
7 until ref not in DB.references
8 return ref

```

3.3 Threat Model

We assume an adversary with unauthenticated network access to the public verification endpoint only (no access to the administrator interface or the underlying database), whose goal is either to (a) locate a reference number that the system will report as valid without having been issued one, in order to fraudulently claim a nonexistent credential, or (b) enumerate the full space of valid references for a given institution/year/type/department combination to harvest personal data disclosed on successful verification. We assume the attacker knows the reference-number format, since it is documented in the system's own scope statement and is, in any case, inferable from a small number of legitimately issued samples.

3.4 Entropy and Collision Analysis of the Baseline Scheme

The suffix is drawn from $16^4 = 65,536$ possible hexadecimal values, i.e., 16 bits of entropy. Table 2 reports the closed-form probability, via the birthday bound, that at least one internal collision is generated over the life of a single (year, type, department) bucket as its record count grows, alongside the mean number of extra retry-loop iterations measured by Monte Carlo simulation (200 trials per bucket size, Python's uuid4 substituted with an equivalent pseudorandom draw over the same 65,536-value space for reproducibility).

Documents per bucket (n)	$P(\geq 1$ internal collision), closed-form	Mean retry-loop iterations (Monte Carlo, 200 trials)
50	1.85%	0.01
100	7.28%	0.08
200	26.21%	0.34
300	49.61%	0.61
500	85.17%	1.69
1000	99.95%	~4 (extrapolated)

Two points follow. First, the retry-on-collision loop does its job: because it re-draws on a detected collision, the system will never issue two documents with the same reference, and the mean cost of doing so stays under two extra iterations even at 1,000 documents per bucket. Second, and separately, the probability that an attacker guesses an existing reference on a given attempt is unrelated to this internal collision-avoidance guarantee; it is set purely by the 65,536-value suffix space. Table 3 quantifies the resulting exposure by estimating the time required to exhaustively enumerate that space against the public verification endpoint at several request rates, in the absence of rate limiting (which the reviewed implementation does not apply).

Attacker request rate	Time to enumerate full 65,536-value suffix space	Practical implication
1 req/s	18.2 hours	Feasible unattended overnight
10 req/s	1.82 hours	Feasible within a work session
50 req/s	21.8 minutes	Trivial with a basic script
100 req/s	10.9 minutes	Trivial; no rate limiting present in the reviewed implementation

3.5 Proposed Hardened Protocol

We propose replacing the random suffix with a keyed message authentication tag computed over the same public fields plus a monotonically increasing per-bucket serial number, using HMAC-SHA256 (Krawczyk et al., 1997) truncated to a configurable number of hexadecimal characters, and a secret key held only by the issuing server. Algorithm 2 specifies both generation and verification.

Algorithm 2: Proposed HMAC-hardened reference-number generation

Input: `matric_number`, `document_type`, `date_issued`, `secret_key K`

```

1 year <- date_issued.split('-')[0]
2 dept <- upper(matric_number[4:7])
3 code <- DOC_TYPE_MAP[document_type]
4 serial <- NEXT_SEQUENCE(year, code, dept) //
monotonic per bucket
5 payload <- year || code || dept || serial
6 tag <- HMAC_SHA256(K, payload)[0:12 hex chars] //
48 bits
7 ref <- "NWU-" + year + "-" + code + "-" + dept + "-" +
serial + "-" + tag
8 return ref

```

Verification(ref, K):

```

1 parse year, code, dept, serial, tag_received from ref
2 tag_expected <- HMAC_SHA256(K,
year||code||dept||serial)[0:12 hex chars]
3 if tag_received == tag_expected AND record exists in
DB: return VALID
4 else: return INVALID

```

Because the tag is keyed, an adversary who knows the format, the serial numbering scheme, and even every previously issued reference gains no advantage in predicting the next valid tag beyond exhaustive search of the truncated HMAC output space (Menezes et al., 1996). Table 4 compares the resulting entropy and brute-force exhaustion time of the baseline against three truncation choices for the proposed scheme.

Scheme	Keyspace	Entropy (bits)	Time to exhaust @ 100 req/s
Baseline (4-hex UUID suffix)	65,536	16	10.9 minutes
Proposed HMAC-32	4.3×10^9	32	~1.4 years
Proposed HMAC-48 (recommended)	2.8×10^{14}	48	~89,000 years
Proposed HMAC-64	1.8×10^{19}	64	~ 5.8×10^9 years

A 48-bit truncation is recommended as the practical default: it keeps the reference number a similar length to the original (twelve hexadecimal characters versus four), remains short enough to transcribe or read aloud, and raises exhaustion time from 11 minutes to an estimated 89,000 years at the same 100-requests-per-second attack rate. We benchmarked the computational overhead of the proposed tag directly: 200,000 HMAC-SHA256 computations completed in 0.645 seconds on commodity hardware, or approximately 3.2 microseconds per operation - three to four orders of magnitude below the system's own measured sub-second verification-query latency, so the hardening carries no perceptible performance cost. We additionally recommend pairing the hardened identifier with basic rate limiting on the verification endpoint (for example, capping requests per source IP per minute) as defence in depth, since entropy alone bounds average-case attacker effort but does not by itself prevent sustained automated probing.

3.6 Comparative Positioning

Table 5 situates the baseline and proposed schemes against manual verification and the QR-code- and blockchain-based approaches reported in the reviewed literature, across infrastructure requirement, offline transcribability by the credential holder, forgery/enumeration resistance, and deployment cost.

System	Extra hardware/printing	Offline transcribable	holder-Forgery/enumeration resistance	Deployment cost
Manual correspondence verification	None	N/A	High (human check) but slow	Low, but high recurring labour cost
QR-code based (Adewale et al., 2020)	QR printing/scanning	No	High, but tied to a physical artefact	Moderate
Blockchain-anchored (Tapscott & Tapscott, 2016)	None	Yes	Very high	High (node/ledger infrastructure)
NWU baseline (this study, as-built)	None	Yes	Low (16-bit, enumerable in minutes)	Very low
NWU hardened, HMAC-48 (proposed)	None	Yes	High (48-bit, ~89,000 yrs @100 req/s)	Very low

Table 5 is necessarily qualitative for the manual, QR-code, and blockchain rows, since those systems were not re-implemented and timed under identical conditions in this study; it is included to situate infrastructure and forgery-resistance trade-offs rather than to claim controlled timing parity across systems. Table 5a supplements it with a narrower, quantitative comparison limited to the two systems and one literature-reported figure for which directly comparable timing evidence is available.

Table 5a. Quantitative verification-time comparison

Verification Method	Typical Verification Time	Basis
Manual correspondence (status quo)	Hours to days	Reported turnaround for institutional written-request verification (Okonkwo & Okafor, 2020)
NWU baseline (as-built)	< 0.5 s (measured, n=100 records)	Directly measured system-testing query latency (Table 6, performance requirement)
NWU hardened (proposed HMAC-48)	< 0.5 s + ~3.2 microseconds	Baseline query latency plus measured HMAC tag-verification cost (Section 3.5)

3.7 Evaluation Strategy

The functional system was evaluated through unit, integration, and system testing, and through a structured user-acceptance exercise with eight participants (four graduating students, two registry staff, two employer/HR representatives) rating the system on a five-point Likert scale after completing predefined tasks. In addition, and specific to this paper's contribution, two security test cases were executed: a simulated full-space enumeration probe against the baseline suffix under a controlled test database, and a verification-forgery attempt against the proposed

HMAC tag using a serial number harvested from a legitimately issued reference but without knowledge of the server key.

The user-acceptance sample (n=8) reflects the population practically reachable within a single-institution, single-semester final-year project - the full set of registry staff and readily available employer contacts willing to participate, plus a convenience sample of graduating students - and is reported here as an exploratory, formative evaluation rather than a statistically representative survey. It is adequate to surface directional stakeholder concerns (Section 4.2 shows it did exactly this for security confidence) but is not powered to support population-level inference, and no statistical significance claims are made from it in this paper. Scaling the evaluation to a survey-representative sample (on the order of 200 respondents stratified across student, staff, and employer sub-populations) is identified as a specific item for the field-trial recommendation in Section 5.0, along with logging individual, respondent-level scores - rather than group means alone - so that the variance within each stakeholder group, and not only its average, can be reported.

4.0 Results and Discussion

4.1 Functional and Security Testing

Table 6 summarises results across all testing levels, including the two security-specific test cases introduced in this study.

Testing Level	Test Cases	Passed	Failed	Pass Rate
Unit testing	16	16	0	100%
Integration testing	7	7	0	100%
System testing	11	11	0	100%
Security testing (enumeration probe, this study)	2	2	0	100%

The enumeration probe against the baseline scheme, run against a seeded 65,536-entry test space, successfully located all valid references within the modelled time bound reported in Table 3, confirming the analytical result empirically rather than only theoretically. The forgery attempt against the HMAC-tagged scheme, using a correct serial number but no access to the server key, failed on every one of 10,000 attempted tags, consistent with the negligible per-guess success probability implied by Table 4.

4.2 User Acceptance

Table 7 reports the group-mean user-acceptance ratings collected for the deployed (baseline) system, with the sample size for each stakeholder group stated in the table caption ($n=8$ overall: 4 students, 2 registry staff, 2 employers). Only group means were recorded during the original evaluation session; individual, respondent-level scores were not retained, so a standard deviation cannot be computed retrospectively from the data available for this study, and none is reported. This is flagged explicitly as a corrected limitation of the present evaluation: a field-trial replication should log every respondent's raw score against each criterion so that both the mean and its dispersion (standard deviation, and ideally a confidence interval appropriate to the resulting sample size) can be reported, rather than the group mean alone.

Within those constraints, confidence in system security was the lowest-rated criterion across all three stakeholder groups (4.3-4.5, against 4.5-4.9 on every other criterion). This is the one dimension on which participant intuition and the entropy analysis in Section 3.4 independently agree: users were, without being shown the underlying mathematics, less confident in precisely the property this paper shows to be measurably weak. We treat this convergence as a suggestive, not statistically established, observation given the sample-size and dispersion limitations just noted; the substantive security claim in this paper rests on the entropy and brute-force analysis of Sections 3.4-3.5, with the UAT result offered as corroborating, formative evidence rather than as independent statistical proof.

4.3 Discussion

These results support the alternative hypothesis that a web-based verification system improves on manual, correspondence-based verification in efficiency and accessibility, while also showing that 'automated' and 'secure' are not synonymous: the baseline system's uniqueness guarantee (Section 3.4) coexists with a guessability weakness that a resourced adversary could exploit in minutes. The proposed HMAC-hardened protocol closes this gap at a computational cost (3.2 microseconds per tag) that is immaterial relative to the system's own sub-second response-time requirement, and without adding printing, scanning, or distributed-ledger infrastructure of the

kind required by the QR-code- and blockchain-based alternatives in Table 5. The wider implication for the applied literature on Nigerian institutional verification systems is that 'unique reference number' claims should be accompanied by an explicit entropy figure and, ideally, a stated request-rate assumption under which brute-force exhaustion becomes impractical; this paper's method (closed-form birthday-bound analysis, Monte Carlo validation, and a measured brute-force feasibility table) is offered as a reusable template for that purpose.

The study's scope should be read with three limitations in mind. First, the hardened protocol was analysed and benchmarked but not deployed for a live multi-user UAT round of its own; the entropy and cost figures reported here are the outcome of controlled simulation and micro-benchmarking rather than field observation. This was a scope decision rather than an oversight: replacing an already-documented reference-number format at NWU requires institutional sign-off from the registry and ICT units that issue and rely on existing references, a migration or dual-running plan for records already created under the baseline format, and coordination that falls outside the single-semester timeline available for an undergraduate final-year project; the hardened scheme is accordingly presented here as an analysed and benchmarked design ready for that follow-on institutional process, not as a claim of field validation. Second, the user-acceptance sample size and the absence of retained respondent-level scores (Section 4.2) limit the evaluation to formative, exploratory status rather than statistically representative measurement. Third, the threat model addresses enumeration and forgery via the public verification endpoint only; it does not cover compromise of the administrator credentials or the server's secret key itself, which would require the key-management and infrastructure-hardening measures noted in the recommendations below.

5.0 Conclusion and Recommendations

This paper has shown that a reference-number scheme's uniqueness guarantee, however rigorously enforced by a retry-on-collision loop, does not by itself establish resistance to guessing, and has quantified that gap for a deployed institutional verification system: a 16-bit suffix space enumerable in under eleven minutes at a modest attack rate. We proposed, specified, and benchmarked a keyed HMAC-SHA256 replacement that raises brute-force exhaustion time by roughly nine orders of magnitude at a per-operation cost of 3.2 microseconds, and showed that it, unlike QR-code- or blockchain-based alternatives, requires no additional printing or ledger infrastructure.

We recommend that Northwest University adopt the 48-bit HMAC-tagged reference scheme with server-side key management (e.g., via an environment-provided secret and

periodic key rotation), pair it with basic per-IP rate limiting on the verification endpoint, and combine both with the previously identified baseline hardening needs - password hashing, CSRF protection, and enforced HTTPS - before any production deployment. We further recommend that future work conduct a live field trial of the hardened protocol with an independent user-acceptance sample, and extend the enumeration-resistance analysis to cover multi-institution deployments where reference numbers may be shared across a federated verification network.

Acknowledgement

The authors acknowledge the Department of Computer Science, Faculty of Computing, Northwest University, Kano, for institutional support during this study.

Authors' Contributions

F.A.F. designed and implemented the baseline system and the proposed hardened protocol, conducted the entropy analysis, simulations, and testing, and drafted the manuscript;

S.B.S. supervised the research and reviewed the manuscript.

Conflict of Interest

The authors declare no conflict of interest.

References

- Adewale, O. S., Boyinbode, O. K., & Salami, H. O. (2020). Development of a web-based certificate verification system using QR code technology. *International Journal of Advanced Computer Science and Applications*, 11(4), 45-52.
- Alsobhi, H. A., Alakhtar, R. A., Ubaid, A., Hussain, O. K., & Hussain, F. K. (2023). Blockchain-based micro-credentialing system in higher education institutions: Systematic literature review. *Knowledge-Based Systems*, 265, 110238.
- Berrios Moya, J. A., Ayoade, J., & Uddin, M. A. (2025). A zero-knowledge proof-enabled blockchain-based academic record verification system. *Sensors*, 25(11), 3450.
- Boonkrong, S. (2024). Design of an academic document forgery detection system. *International Journal of Information Technology*.
- Danish, M., Sobrados, E., Kaushik, P., Acharya, B., Saad, M., Mueen, A., Rahaman, S., & Anwar, A. (2026). Private links, public leaks: Consequences of frictionless user experience on the security and privacy posture of SMS-delivered URLs (arXiv:2601.09232). arXiv.
- Grinberg, M. (2018). *Flask web development: Developing web applications with Python* (2nd ed.). O'Reilly Media.
- Heeks, R. (2006). *Implementing and managing eGovernment: An international text*. SAGE Publications.
- Hipp, D. R. (2023). SQLite: A self-contained, high-reliability, embedded, full-featured, public-domain, SQL database engine. <https://www.sqlite.org>
- Krawczyk, H., Bellare, M., & Canetti, R. (1997). HMAC: Keyed-hashing for message authentication (RFC 2104). Internet Engineering Task Force.
- Menezes, A. J., van Oorschot, P. C., & Vanstone, S. A. (1996). *Handbook of applied cryptography*. CRC Press.
- MITRE. (2024). CWE-331: Insufficient entropy. Common Weakness Enumeration. <https://cwe.mitre.org/data/definitions/331.html>
- National Information Technology Development Agency (NITDA). (2021). Report on digital identity and credential management in Nigeria. Federal Ministry of Communications and Digital Economy.
- Okonkwo, C. W., & Okafor, A. C. (2020). Academic credential fraud in Nigerian universities: Survey findings and systemic analysis. *Journal of Education Administration and Policy*, 6(2), 112-128.
- Olayemi, O. J., Tijani, A. H., & Lawal, B. A. (2019). A review of digital document management and authentication systems in higher education. *Journal of Information Systems and Technology Management*, 16, e201916010.
- OWASP. (2021). OWASP top ten: The ten most critical web application security risks (2021 ed.). Open Web Application Security Project. <https://owasp.org/www-project-top-ten/>
- Pressman, R. S., & Maxim, B. R. (2015). *Software engineering: A practitioner's approach* (8th ed.). McGraw-Hill Education.
- Ramić, Š. B., Prazina, I., Pozderac, D., Mulahasanović, R. T., & Mrdović, S. (2024). BLS-MT-ZKP: A novel approach to selective disclosure of claims from digital credentials. *IEEE Access*, 12, 192062-192078.
- Royce, W. W. (1970). Managing the development of large software systems: Concepts and techniques. *Proceedings of IEEE WESCON*, 26, 1-9.
- Sheffer, Y., Hardt, D., & Jones, M. B. (2020). JSON Web Token best current practices (RFC 8725). Internet Engineering Task Force.
- Tapscott, D., & Tapscott, A. (2016). *Blockchain revolution: How the technology behind Bitcoin is changing money, business, and the world*. Portfolio/Penguin.
- Xu, Y. (2024). Development of blockchain-based academic credential verification system. *Open Access Library Journal*, 11(9), 1-20.