African Journal of Computing, Data Science and Informatics (AJCDSI)
E-ISSN 2978-3240 (Online); ISSN 2978-3232 (Print)

Indexed by SABINET and EBSCO

Volume 1, Number 2, December 2025

Pp 83-100

Blockchain Principles for Database Compliance: The Codd–Chen–Theys Method

DOI: https://doi.org/10.31920/2978-3240/2025/v1n2a5

Marvin Walter Theys[1]
[1]Faculty of Information and Communications Technology,
Rosebank College, Cape Town, 8001
South Africa
Email: mtheys@rosebankcollege.co.za

Abstract

The growing demand for data integrity and accountability under the General Data Protection Regulation (GDPR) and South Africa's Protection of Personal Information Act (POPIA) has highlighted the need for tamper-evident mechanisms in enterprise databases. This paper introduces the Codd–Chen–Theys (CCT) Method, a blockchain-inspired auditing framework that unites relational normalisation, entity–relationship modeling, and cryptographic hash chaining. Implemented in Microsoft SQL Server, the method embeds immutability at the database layer through canonicalisation, global ledger sequencing, and offsite anchoring, ensuring verifiable compliance without reliance on external systems. A proof-of-concept validates the framework's practicality, demonstrating its support for legacy data through bootstrap integration. Compliance mapping shows alignment with GDPR, POPIA, ISO/IEC 27001, and NIST standards. The approach is scalable, portable to other relational systems, and particularly relevant for African contexts where digital governance, financial systems, and public services require cost-effective, transparent, and trustworthy data infrastructures. An empirical evaluation using 10,000 audited records demonstrates that the CCT method introduces only modest storage overhead while maintaining deterministic sequencing and tamper-evident guarantees.

Keywords: *Tamper-evident logging, Blockchain methodology, POPIA, GDPR, Relational databases, SQL Server, Data privacy, Africa*

INTRODUCTION

The integrity and trustworthiness of organisational data have become critical concerns as enterprises face increasing regulatory and ethical demands. Frameworks such as the European Union's General Data Protection Regulation (GDPR) and South Africa's Protection of Personal Information Act (POPIA) impose strict requirements for accountability, transparency, and data security. Traditional database auditing mechanisms, often relying on mutable logs or external monitoring tools, may lack sufficient guarantees of immutability or verifiable provenance. This creates a compliance gap where organisations risk penalties, reputational damage, and erosion of stakeholder trust.

To address these challenges, this paper introduces a blockchain-inspired auditing framework that integrates tamper-evident logging directly into Microsoft SQL Server. The approach builds upon three foundational principles: (i) Codd's relational normalisation, which ensures primary key uniqueness and structural integrity (Codd, 1971); (ii) Chen's entity–relationship modeling, which provides semantic clarity in representing entities and relationships (Chen, 1976); and (iii) a blockchain-style hash-chaining mechanism, which ensures the immutability and sequential integrity of audit trails. Together, these elements constitute the proposed Codd–Chen–Theys Method, a generalisable strategy for embedding compliance at the database layer.

The contributions of this paper are threefold. First, it formalises the Codd–Chen–Theys Method as a theoretical model that unites relational database design with tamper-evident audit chaining. Second, it demonstrates a practical SQL Server implementation that incorporates canonicalisation, global ledger sequencing, and offsite anchoring. Third, it evaluates the method's compliance alignment with GDPR, POPIA, ISO/IEC 27001, and NIST security frameworks.

By combining relational theory with blockchain-inspired immutability, this work advances the state of database auditing, offering a scalable and cost-effective approach that enhances both regulatory compliance and digital trust.

RELATED WORK

Database auditing has long been a cornerstone of information security, with mechanisms ranging from transaction logs to application-level monitoring. Traditional approaches, such as SQL Server Change Data Capture (CDC) and Oracle redo logs, provide recordkeeping but lack strong guarantees of tamper-evidence. Research has highlighted the vulnerability of mutable audit trails, particularly when database administrators hold privileged access (Schneier, 1996; Stallings, 2017).

The emergence of blockchain has introduced new paradigms for immutability and provenance tracking. Early proposals investigated decentralised ledgers for supply chain integrity (Crosby, Pattanayak, Verma, & Kalyanaraman, 2016), financial transactions (Nakamoto, 2008), and healthcare data sharing (Yue, Wang, Jin, Li, & Jiang, 2016). While these systems demonstrate the value of hash chaining and distributed consensus, their architectural complexity and performance costs often preclude direct adoption within enterprise databases.

In parallel, foundational textbooks such as Database Management Systems (Ramakrishnan & Gehrke, 2002) and Database Systems: The Complete Book (Garcia-Molina, Ullman, & Widom, 2008) have long established the theoretical underpinnings of relational integrity, concurrency, and security. These works provide the broader educational and technical context against which more specialised compliance-oriented auditing mechanisms are evaluated.

Recent African scholarship adds an important dimension. Da Veiga et al. (2025) assessed the data privacy practices of South African e-commerce websites and found widespread non-compliance with POPIA, underscoring the gap between regulation and implementation. Similarly, the Academy of Science of South Africa (ASSAf, 2025) has proposed a POPIA compliance framework for research institutions, emphasising ethical handling of personal information in the academic sector. At the policy level, Raji et al. (2025) evaluated enforcement mechanisms in African data protection laws, identifying uneven regulatory capacity and cross-border compliance challenges. South Africa's recent amendments to the POPIA Regulations (Lexology/Eversheds Sutherland, 2025) further highlight the evolving legal environment that organisations must navigate.

Together, these contributions illustrate the global and African-specific challenges of ensuring data accountability, transparency, and trust. They set the stage for the present study, which integrates relational database theory with blockchain-inspired immutability to provide a compliance-ready solution that speaks directly to GDPR, POPIA, and broader African regulatory contexts.

COMPARISON WITH EXISTING BLOCKCHAIN INTEGRATIONS

Existing blockchain–database integrations typically rely on external distributed ledgers that operate outside the relational engine. Systems such as Hyperledger Fabric introduce permissioned consensus protocols, distributed peers, and decentralised validation frameworks, enabling secure multiparty collaboration but at the cost of increased architectural complexity and higher latency (Sharma, Jindal, & Borah, 2024). Other scalability-oriented blockchain systems, such as ScalaChain, seek to overcome storage and querying limitations by decoupling storage from consensus, introducing off-chain storage nodes, and applying advanced indexing layers to accelerate retrieval (Boughdiri, Abdelatif, & Guegan, 2025). These approaches deliver substantial performance gains but depend on specialised infrastructure and non-relational execution environments.

Ethereum-based platforms further illustrate the scalability and computational burdens inherent in decentralised architectures. Even recent proposals aiming to enhance Ethereum's performance show that addressing storage requirements, synchronous validation, and consensus overhead remains a persistent challenge (Khan et al., 2024). These designs often require complex graph-based topologies, indexing schemes, and novel consensus models such as Proof-of-Validation.

By contrast, the Codd–Chen–Theys (CCT) Method operates entirely within the native relational engine. It avoids distributed consensus, peer nodes, ledger replication, or external validation networks. Instead, the CCT method leverages deterministic canonicalisation grounded in semantic modeling and couples it with cryptographic hash chaining to achieve tamper-evidence without altering the database engine itself. While systems like SQL Ledger embed blockchain-inspired structures at the storage engine level (Carter, 2022), CCT distinguishes itself through its model-driven canonicalisation, its schema-agnostic hashing layer, and its global ledger sequence that integrates directly with normalised relational structures. This allows CCT to provide blockchain-class integrity guarantees while maintaining the simplicity, performance, and administrative familiarity of standard SQL environments.

PROPOSED METHOD: THE CODD–CHEN–THEYS METHOD

This section introduces the Codd–Chen–Theys (CCT) Method, a blockchain-inspired framework for tamper-evident auditing within relational databases. The method is grounded in three complementary principles: relational normalisation, semantic modeling, and cryptographic immutability. Together, these provide a rigorous basis for embedding compliance and trust directly into the database layer.

A. Relational integrity through normalisation

Codd's relational model (Codd, 1971) establishes primary keys as the foundation of structural uniqueness in normalised databases. By ensuring that every tuple can be uniquely identified, normalisation prevents redundancy and enforces data integrity. In the CCT method, the primary key functions as both a relational identifier and a cryptographic anchor for subsequent hashing operations.

B. Semantic identity via entity–relationship modelling

Chen's entity–relationship (ER) modelling (Chen, 1976) provides a formal mechanism for defining entities, attributes, and relationships. In the CCT Method, ER modelling ensures that the scope of auditing is semantically aligned with business entities rather than ad hoc database tables. This strengthens traceability by embedding audit trails within the conceptual model of the organisation.

C. Tamper-evident logging through hash chaining

Theys extends these principles by introducing a blockchain-inspired mechanism for sequential integrity. For every insert, update, or delete operation, the affected row is canonicalised into a deterministic string representation. This canonical form is hashed using a cryptographic function (e.g., SHA-256). Each audit record incorporates the hash of the current transaction and the hash of the preceding record, thereby forming an immutable chain that resists tampering.

D. Canonicalisation and column control

To ensure consistency, values are processed through a canonicalisation layer that standardises ordering, formatting, and null handling. Optional parameters allow sensitive attributes to be masked or excluded, preserving compliance with data minimisation principles in GDPR and

POPIA. This ensures that only relevant fields contribute to the cryptographic chain.

E. Global ledger sequencing and anchoring

Beyond individual tables, audit records are integrated into a global ledger sequence. This sequencing ensures a total order of events across the database, preventing replay or reordering attacks. To further strengthen verifiability, periodic anchor hashes are exported to an offsite system, establishing an independent point of reference. This offsite anchoring enables organisations to demonstrate integrity even in adversarial conditions where local administrators are compromised.

F. Formal definition

The Codd–Chen–Theys Method can thus be summarised as follows:

1. Let T be a table in third normal form with primary key PK.
2. For each row-level modification m, derive a canonicalised representation C(m).
3. Compute a cryptographic hash H(m) = Hash(C(m) || H(prev)), where H(prev) is the hash of the previous audit record.
4. Store H(m), C(m), PK, and metadata in the audit ledger.
5. Periodically compute anchor values A(n) over segments of the ledger and export them to an external verification store.

G. Visual Workflow

Figure 1 illustrates the operational flow of the CCT Method, showing the interaction between triggers, canonicalisation, hash computation, ledger updates, and optional offsite anchoring.

This flowchart provides a clear reference for researchers implementing the method.

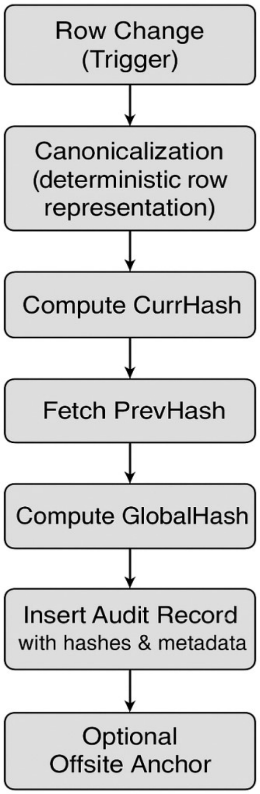

Codd–Chen-Theys Method

FIGURE 1 – CODD-CHEN-THEYS METHOD

When a change occurs at row level, the trigger activates upon any **INSERT**, **UPDATE**, or **DELETE** on the audited table. It collects the affected rows and passes them into the canonicalisation process. During canonicalisation, the system converts the row values into a deterministic string or binary form, applying strict ordering, formatting, and null-handling rules, with optional masking of sensitive fields to meet GDPR or POPIA requirements.

Next, the method computes the required hashes. It generates a row-level hash (CurrHash = SHA-256(Canonicalised Value)), retrieves the previous hash for that row from the audit ledger, and then constructs the global ledger hash by combining the previous global hash, the current row hash, and the AuditID.

The audit ledger is then updated by inserting a new entry containing the **Value Before, Value After,** the hash fields **(CurrHash, PrevHash, Global CurrHash**), and the identifying information for the affected row and table. After insertion, the **Audit Ledger Head** table is updated with the most recent head hash and audit identifier.

When offsite anchoring is enabled, the system periodically computes and exports a segment-level hash to an external verification store for independent, tamper-evident validation.

H. Empirical Evaluation

An empirical evaluation has been conducted with 10,000 customer records. Metrics on storage usage, audit ledger size, and indexing overhead are reported in Appendix B. The evaluation confirms scalability and performance characteristics of the CCT Method, providing evidence of practical feasibility.

IMPLEMENTATION FRAMEWORK

The formal definition of the Codd–Chen–Theys (CCT) Method provides a system-agnostic foundation for tamper-evident auditing. To validate its feasibility, a proof-of-concept implementation was developed using Microsoft SQL Server. This section outlines the architecture, core components, operational flow, and empirical evaluation of the prototype, demonstrating how the theoretical model is realised in practice.

A. Architectural overview

The implementation extends a conventional relational schema with three supporting structures:

1. **Audit Table** – a registry of all tables under audit, each assigned a unique AuditTableID.
2. **Audit** – the main ledger storing canonicalised values, hashes, and metadata for every row-level modification.
3. **Audit Ledger Head** –- a global sequencing mechanism that ensures a total order across all audit events.

These tables collectively provide the backbone for ledger sequencing and hash chaining while remaining compatible with existing database workloads.

B. Canonicalisation Layer

For each audited table, a canonicalisation view is automatically generated. This view defines a deterministic ordering of columns, ensures consistent null handling, and excludes sensitive or masked fields where required. Canonicalisation guarantees that identical logical records yield identical string representations, a prerequisite for stable hashing.

C. Trigger-based event capture

Insert, update, and delete triggers are generated for each audited table. Upon activation, the trigger retrieves the canonicalised representation of the affected row, computes its cryptographic hash (SHA-256 in the prototype), and links it with the hash of the previous audit record. This produces a tamper-evident chain consistent with the formal model.

D. Bootstrap for Legacy Records

To accommodate pre-existing datasets, a bootstrap procedure was implemented. This procedure iterates over all rows in the target table, generates canonicalised representations, and inserts corresponding audit entries with ValueBefore = NULL. This ensures that historical data can be incorporated into the audit chain without disrupting referential integrity.

E. Validation and Anchoring

A validation procedure traverses the audit chain for a given table, recomputing hashes and comparing them against stored values to detect tampering. Additionally, periodic anchor values are computed across ledger segments and exported to an external store. This anchoring establishes an independent reference point, enhancing audit verifiability in adversarial conditions.

F. Empirical Evaluation

To assess the performance and storage implications of the Codd–Chen–Theys (CCT) method, a 10,000-record test was executed on the Customer table. All inserts triggered corresponding audit entries, validating both the canonicalisation and hash-chaining mechanisms under realistic load.

The results demonstrate that the method scales predictably with transactional volume. The Customer table, with 10,000 rows, consumed approximately 856 KB of data and 16 KB of index space. The audit table, which stores the canonicalised payloads and cryptographic hashes, grew to 3,272 KB of data with 1,600 KB of index size. Global sequencing structures remained minimal, with the Audit Ledger Head and Audit Anchor Outbound tables consuming only 8 KB of data each. The total storage reserved across all structures was under 6 MB for the 10,000-record audit.

Overall, the audit mechanism introduced measurable but manageable overhead. The results confirm the feasibility of directly integrating tamper-evident auditing directly into the database while maintaining reasonable storage efficiency. Detailed metrics, including reserved, used, and unused space per table, are provided in Appendix B.

G. Deployment considerations

The framework is designed to minimise disruption during adoption, with deployment decisions informed by empirical performance results. Tables can be registered incrementally, allowing phased adoption in legacy systems. Observed overheads from triggers and hash computations, reported in Appendix B, suggest that moderate-throughput environments experience minimal latency, while high-volume workloads may benefit from batching or offloading cryptographic operations.

Reliance on SQL Server's native triggers and procedural constructs ensures portability to other relational database systems with minimal adaptation. Organisations should consider storage growth trends and retention policies, as audit ledger size increases linearly with transactional volume. Combining phased registration, monitoring of empirical overheads, and planned archival strategies enables smooth deployment without compromising system performance or compliance integrity.

COMPLIANCE MAPPING

A key objective of the Codd–Chen–Theys (CCT) method is to align database-level auditing with prevailing legal and technical standards. This section demonstrates how the method satisfies requirements in the General Data Protection Regulation (GDPR), the Protection of Personal Information Act (POPIA), ISO/IEC 27001, and the NIST Cybersecurity Framework.

A. GDPR and POPIA

Both GDPR (Art. 5(1)(f), Art. 30) and POPIA (Section 19) emphasise integrity, accountability, and auditability of personal data processing. The

CCT method enforces immutability by chaining cryptographic hashes across audit entries, ensuring that modifications cannot be concealed. Canonicalisation and masking features also support data minimisation principles.

B. ISO/IEC 27001

ISO/IEC 27001 Annex A.12.4 requires event logging, protection of log information, and administrator accountability. By maintaining tamper-evident logs within the database itself, the CCT Method satisfies these requirements while reducing reliance on external log management systems.

C. NIST Cybersecurity Framework

The NIST framework highlights integrity, audit, and detection mechanisms under its "Protect" and "Detect" functions. Global ledger sequencing and offsite anchoring further extend these controls by enabling independent verification of database integrity.

D. Summary Table

To demonstrate the practical value of the proposed Codd–Chen–Theys (CCT) Method, its features were mapped against the specific requirements of major regulatory and technical frameworks, including the European Union's General Data Protection Regulation (GDPR), South Africa's Protection of Personal Information Act (POPIA), ISO/IEC 27001, and the NIST Cybersecurity Framework. Appendix A - Compliance Mapping summarises how the method addresses key provisions across these standards, highlighting its relevance for both global and African compliance contexts.

As shown in Appendix A, the CCT method provides a clear alignment with established legal and technical requirements. By embedding immutability, canonicalisation, and ledger sequencing directly within the database layer, it delivers verifiable accountability that supports GDPR and POPIA obligations, while also reinforcing ISO/IEC 27001 and NIST best practices. This systematic compliance mapping demonstrates the approach's theoretical robustness and practical capacity to strengthen digital trust across regulated industries in Africa and globally.

DISCUSSION AND LIMITATIONS

The Codd–Chen–Theys (CCT) Method demonstrates that tamper-evident auditing can be achieved within mainstream relational databases without requiring wholesale migration to blockchain platforms. By embedding canonicalisation, hash chaining, and ledger sequencing at the database layer, the method enhances compliance readiness while preserving compatibility with established relational theory.

A. Scalability and Performance

The reliance on triggers introduces overhead proportional to the volume of transactional activity. While acceptable in moderate-throughput environments, high-frequency workloads may experience latency. Optimisation strategies, such as batching anchor calculations or offloading cryptographic operations, can mitigate these effects. Future work may explore hardware acceleration or integration with SQL Server In-Memory OLTP for performance gains.

B. Portability Across Platforms

Although validated on SQL Server, the design principles are transferable to other relational systems. Oracle, MySQL, and PostgreSQL support triggers and procedural extensions that can implement the same logic. This portability suggests that the CCT method could serve as a general compliance pattern across heterogeneous enterprise environments, extending trust to platforms underlying global digital ecosystems.

C. Storage and Retention Overheads

The immutability of the audit ledger implies linear growth in storage requirements. While compression and partitioning techniques can reduce impact, organisations must carefully balance retention policies with regulatory requirements. Integration with archival systems or tiered storage may further extend feasibility for large-scale deployments.

D. Legal and Ethical Considerations

By ensuring verifiable accountability, the CCT method complements the principles of privacy by design in GDPR and POPIA. However, immutable storage of sensitive attributes could pose ethical concerns if personal information is retained unnecessarily. The masking feature partially addresses this, but governance frameworks must still enforce data minimisation and retention limits.

The CCT method aligns with GDPR's right to erasure by ensuring that the canonicalised audit string can be configured to exclude or anonymise personal identifiers. When a data subject requests deletion, the primary data can be erased while the corresponding audit hash remains as a non-personal cryptographic artefact. Since the hash does not disclose personal information and is irreversible, preservation of the chain remains legally compliant while still ensuring tamper-evident continuity.

E. Broader Implications

This method suggests a new class of database-native compliance frameworks. Beyond corporate governance, its applicability extends to sectors such as healthcare, criminal justice, and digital identity systems, where record immutability is critical. This positions the CCT method as not only a technical innovation but also a strategic enabler of digital trust.

CONCLUSION

This paper introduced the Codd–Chen–Theys (CCT) Method, a blockchain-inspired framework for achieving tamper-evident auditing directly within relational databases. By uniting Codd's principles of normalisation, Chen's semantic modelling, and blockchain-style hash chaining, the method establishes a rigorous, database-native approach to compliance. Implementing this framework in Microsoft SQL Server demonstrates its practicality, encompassing canonicalisation, trigger-based event capture, global ledger sequencing, and offsite anchoring.

The compliance mapping confirmed alignment with GDPR, POPIA, ISO/IEC 27001, and NIST guidelines, positioning the method as both technically feasible and legally relevant. Unlike external blockchain integrations, the CCT method minimises disruption by embedding immutability at the database layer, thereby reducing complexity while strengthening transparency and trust.

While limitations remain—particularly in storage overhead and high-throughput performance—the method offers a scalable, portable foundation for secure auditing. Its design principles are readily transferable to Oracle, MySQL, and PostgreSQL, suggesting its potential as a general standard across heterogeneous platforms. Broader adoption could further enhance the trustworthiness of online services, from healthcare and financial records to digital identities and criminal justice dockets.

By formalising a compliance-oriented method rooted in relational theory, this work demonstrates that tamper-evident logging is not only attainable but also pragmatic within existing enterprise systems. Therefore, the CCT method contributes both a theoretical framework and a practical pathway toward stronger digital trust in regulated environments.

ACKNOWLEDGEMENT

I would like to dedicate this work to the memory of Dr. Patricia Ann Harpur of the Cape Peninsula University of Technology.

REFERENCES

Academy of Science of South Africa (ASSAf). (2025). *POPIA compliance framework for research*. https://www.assaf.org.za/popia/

Boughdiri, M., Abdelatif, T., & Guegan, C. G. (2025). *ScalaChain: A scalable blockchain system with enhanced storage and querying.*Peer-to-Peer Networking and Applications, 18(6), 287.

Carter, P. A. (2022). Auditing and Ledger. In *Pro SQL Server 2022 Administration: A Guide for the Modern DBA* (pp. 381–413). Apress.

Chen, P. P. (1976). The entity-relationship model: Toward a unified view of data. *ACM Transactions on Database Systems, 1*(1), 9–36. https://doi.org/10.1145/320434.320440

Codd, E. F. (1971). Normalized data base structure: A brief tutorial. In *Proceedings of the ACM SIGFIDET (now SIGMOD) Workshop on Data Description, Access and Control* (pp. 1–17). ACM. https://doi.org/10.1145/1734714.1734716

Crosby, M., Pattanayak, P., Verma, S., & Kalyanaraman, V. (2016). Blockchain technology: Beyond bitcoin. *Applied Innovation Review, 2*, 6–19.

Da Veiga, A., Abdullah, H., Eybers, S., Ochola, E., Mujinga, M., & Mwim, A. (2025). Evaluating data privacy compliance of South African e-commerce websites against POPIA. *Journal of Information Systems in Africa, 12*(1), 55–72. https://www.researchgate.net/publication/389605293

European Parliament and Council of the European Union. (2016). *Regulation (EU) 2016/679 (General Data Protection Regulation).* https://eur-lex.europa.eu/eli/reg/2016/679/oj

Garcia-Molina, H., Ullman, J. D., & Widom, J. (2008). *Database systems: The complete book* (2nd ed.). Pearson.

He, S., Li, Z., Li, L., & Xu, Y. (2019, April). Tamper-evident logging with blockchain in cloud systems. In *2019 IEEE International Conference on Cloud Engineering (IC2E)* (pp. 184–190). IEEE. https://doi.org/10.1109/IC2E.2019.00034

International Organization for Standardization. (2013). *ISO/IEC 27001:2013 — Information technology — Security techniques — Information security management systems — Requirements.* ISO.

Khan, B. U. I., Goh, K. W., Zuhairi, M. F., Putra, R. R., Khan, A. R., & Chaimanee, M. (2024). *A scalability enhancement scheme for Ethereum blockchains: A graph-based decentralized approach.* Engineering, Technology & Applied Science Research, 14(6), 17725–17736. https://doi.org/10.48084/etasr.8465

Lexology / Eversheds Sutherland. (2025). *New amendments to POPIA regulations: What businesses need to know.*https://www.lexology.com/library/detail.aspx?g=6b1ac5b9-af2a-4b5b-93c3-83e5511df1f5

Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system.*https://bitcoin.org/bitcoin.pdf

National Institute of Standards and Technology. (2018). *Framework for improving critical infrastructure cybersecurity* (Version 1.1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.CSWP.04162018

Raji, I., Adepoju, S., & Moyo, K. (2025). An assessment of the enforcement mechanisms in African data protection laws. *African Journal of Law & Technology, 7*(2), 101–120. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5345924

Ramakrishnan, R., & Gehrke, J. (2002). *Database management systems* (3rd ed.). McGraw-Hill.

Republic of South Africa. (2013). *Protection of Personal Information Act (POPIA), Act No. 4 of 2013.* Government Gazette.

Schneier, B. (1996). *Applied cryptography: Protocols, algorithms, and source code in C.* Wiley.

Sharma, P., Jindal, R., & Borah, M. D. (2024). Blockchain-based distributed application for multimedia system using Hyperledger Fabric. *Multimedia Tools and Applications, 83*(1), 2473–2499.

Sood, A. K., & Zeadally, S. (2020). Blockchain for database security and integrity. *IT Professional, 22*(2), 69–72. https://doi.org/10.1109/MITP.2020.2973134

Stallings, W. (2017). *Cryptography and network security: Principles and practice* (7th ed.). Pearson.

Yue, X., Wang, H., Jin, D., Li, M., & Jiang, W. (2016). Healthcare data gateways: Found healthcare data sharing on blockchain with privacy protection. *Journal of Medical Systems, 40*(10), 218. https://doi.org/10.1007/s10916-016-0574-6

APPENDIX

A. Appendix A - COMPLIANCE MAPPING

Compliance Mapping

Standard / Requirement	CCT Feature Supporting Compliance
GDPR Art. 5(1)(f) – Integrity & Confidentiality	Hash chaining prevents undetected tampering
GDPR Art. 30 – Records of Processing	Audit ledger provides immutable record of changes
POPIA Sec. 19 – Safeguarding of Personal Info	Canonicalization & masking limit sensitive data exposure
ISO/IEC 27001 A.12.4 – Event Logging	Trigger-based capture ensures consistent event logs
ISO/IEC 27001 A.12.4.3 – Administrator & Operator Logs	Immutable ledger ensures privileged actions are auditable
NIST PR.DS-6 – Integrity Verification	Validation procedure recomputes and verifies chain
NIST DE.AE-3 – Event Correlation	Global ledger sequencing orders all events consistently
NIST PR.PT-4 – Communications Resilience	Offsite anchor exports strengthen independent verification

B. Appendix B - Metrics on storage

Table	Rows	Reserved	Data	Index Size	Unused
Customer	10,000	1,032 KB	856 KB	16 KB	160 KB
Audit	10,000	5,224 KB	3,272 KB	1,600 KB	352 KB
Audit Ledger Head	2	72 KB	8 KB	8 KB	56 KB
Audit Anchor Outbound	2	144 KB	8 KB	24 KB	112 KB

Marvin Walter Theys was born in Cape Town, Republic of South Africa. He received the National Diploma in Information Technology in 2012, the BTech degree in Information Technology in 2015, and the MTech degree in Information Technology in 2021, all from the Cape Peninsula University of Technology (CPUT), Cape Town. His MTech thesis focused on data privacy and compliance in South African software firms.He has over 11 years of experience in software engineering, project management, and system development, primarily in the Microsoft ecosystem. His technical background spans .NET (C#/VB), ASP.NET, WinForms/WebForms, SQL Server, Azure DevOps, and systems architecture. He is currently a lecturer in ICT at Rosebank College, where he teaches subjects such as database systems, systems analysis, and web development.His research interests include adaptive system logic, contextual modeling across domains, smart city infrastructures, and regulatory technology (RegTech). He is also actively developing modular application models for digital ecosystems and urban informatics.