Vol. 3, No. 2, pp.: 48-59, 2025

Tamper-Evident Data Architectures for Cross-Sector Digital Trust: A Comparative Framework for Traceability, Interoperability, and System Resilience

Zhiyuan Mei [1], Kaiwen Lu [2,*], Ruoxi Han [3]

[1] School of Computer Science and Technology, Hangzhou Dianzi University, Hangzhou 310018, China

[2] School of Management Science and Engineering, Zhejiang Gongshang University, Hangzhou 310018, China

[3] School of Information Engineering, Nanjing University of Finance and Economics, Nanjing 210023, China

** kaiwen.lu@zjgsu.edu.cn*

Article Information	
Received	18 January 2025
Accepted	29 May 2025
DOI	https://doi.org/10.63646/datamind.2025.030205

Abstract

Digital trust is increasingly produced not by institutions alone but by the data architectures that record, link, and expose evidence of what happened. This article develops a comparative framework for tamper-evident data architectures—systems that make unauthorized modification of recorded data detectable—and evaluates their fitness for cross-sector digital trust. Rather than treating distributed-ledger technology as a single artefact, we distinguish six recurring architectural patterns, ranging from public permissionless ledgers to centralized hash-chained logs, and assess them along seven methodological dimensions: traceability granularity, interoperability, system resilience, throughput, privacy, governance clarity, and cost-efficiency. Using a structured comparative coding design, we score each pattern and analyse the resulting matrix descriptively and through scenario-weighted suitability ranking for three representative settings: supply-chain traceability, healthcare data exchange, and energy peer-to-peer trading. The analysis shows that no single architecture dominates; fitness is contingent on which trust property a sector prioritizes. Permissioned and hybrid ledgers offer the most balanced profiles for regulated data exchange, public ledgers maximize resilience and auditability at the cost of throughput and privacy, and centralized tamper-evident logs remain competitive where governance and efficiency outweigh decentralization. We argue that traceability, interoperability, and resilience should be treated as explicit, weighted design objectives rather than assumed by-products of a chosen platform, and we offer a selection framework that aligns architectural choice with sector-specific trust requirements. The contribution is conceptual and methodological: a transparent, reproducible basis for choosing tamper-evident architectures before system construction begins.

Keywords: ***Tamper-evident data; digital trust; distributed ledger technology; data provenance; traceability; interoperability; system resilience; cross-sector data sharing***

1. Introduction

Trust in digital systems has historically been delegated to institutions. Banks vouched for balances, registries vouched for ownership, and certifying bodies vouched for credentials. A quieter but more consequential shift is now under way: trust is increasingly produced by the data architecture itself, by the way records are captured, cryptographically linked, replicated, and made independently verifiable. When an architecture can demonstrate that a record has not been altered since it was written, parties who do not trust one another can still agree on a shared version of events. This property, rather than decentralization for its own sake, is what gives distributed ledgers and related designs their practical value across sectors (Lu, 2019b; Lu, 2022).

The literature, however, tends to use the language of "immutability" loosely, as if writing data to a ledger made it permanent in some absolute sense. A more precise and more useful notion is tamper-evidence: the guarantee that any unauthorized modification of stored data becomes detectable, because it would invalidate a chain of cryptographic commitments that many independent parties can recompute. Tamper-evidence does not make data physically unchangeable; it makes silent change infeasible. This distinction matters for design, because it clarifies what an architecture actually provides—verifiable detection of tampering—and what it leaves to governance, such as how to correct an erroneous but validly recorded entry (Zheng & Lu, 2022; Lu, 2018).

A second source of confusion is the treatment of distributed-ledger technology as a single artefact. In practice, the design space is wide. Public permissionless ledgers, permissioned consortium ledgers, hybrid public–private designs, directed-acyclic-graph ledgers, layer-two and sidechain-anchored systems, and even centralized logs that publish periodic cryptographic commitments all deliver tamper-evidence, but they differ sharply in throughput, privacy, governance, and the ease with which they interconnect. Choosing among them as if they were interchangeable back ends is a common and costly mistake, because each pattern embeds assumptions that constrain what kinds of trust claims a system can credibly make (Xu et al., 2021; Chen et al., 2024).

The practical consequence is that organizations frequently select an architecture for reasons of familiarity, vendor availability, or momentum, rather than because its structure matches the trust property the application most needs. A food exporter that primarily needs fine-grained traceability has different requirements from a hospital network that primarily needs confidential, governed data exchange, which in turn differs from an energy microgrid that primarily needs high-throughput settlement among many small participants. Treating these as the same "blockchain problem" obscures the trade-offs that ultimately determine whether a deployment succeeds (Agbo et al., 2019; Leng et al., 2022).

This article responds to that gap by developing a comparative framework rather than proposing yet another platform. We pursue three questions. First, how do the principal tamper-evident architectural patterns differ when assessed along the methodological dimensions that matter for cross-sector trust—traceability, interoperability, resilience, throughput, privacy, governance, and cost? Second, do these patterns separate into coherent families with distinct strengths? Third, which patterns are best matched to representative cross-sector scenarios once the analytical objective is made explicit and the dimensions are weighted accordingly?

The stakes of getting this choice wrong are substantial. When a tamper-evident system is mismatched to its task, organizations may over-engineer trust where a simpler audited database would suffice, paying in latency, cost, and operational fragility; or they may under-engineer it, adopting a lightweight log where adversarial, multi-party conditions demand open replication, and discovering the gap only after a dispute or breach. Both failure modes are expensive, and both are avoidable if the trust requirement is articulated before the platform is selected. Because

records committed to these systems are designed to persist and to be verified by others, early architectural mistakes are unusually difficult to reverse, which raises the value of a disciplined, comparative selection method at the design stage.

To answer these questions we adopt a structured comparative coding design inspired by methodological benchmarking in adjacent fields. We code six architectural patterns against seven dimensions on a comparative one-to-five scale, then analyse the matrix descriptively and through scenario-weighted suitability scoring for supply-chain traceability, healthcare data exchange, and energy peer-to-peer trading. The exercise does not attempt to measure the performance of any specific implementation; it formalizes the qualitative reasoning that practitioners already perform implicitly, so that architectural selection becomes explicit, contestable, and reproducible. The remainder of the paper is organized as follows. Section 2 sets out the conceptual foundations of tamper-evidence and the layered architecture we analyse. Section 3 defines the comparative framework, its patterns, and its dimensions. Section 4 reports the descriptive and scenario-based analysis. Section 5 discusses implications, convergence with adjacent technologies, and limitations, and Section 6 concludes.

2. Conceptual Foundations of Tamper-Evident Data Architectures

Tamper-evidence rests on a small number of cryptographic primitives whose composition produces a powerful systemic property. The first is the cryptographic hash function, a deterministic mapping from arbitrary input to a fixed-length digest for which finding two inputs with the same digest is computationally infeasible. The second is the hash chain, in which each record stores the digest of its predecessor, so that altering any earlier record changes its digest and breaks every subsequent link. The third is the Merkle tree, which summarizes a large set of records under a single root digest, allowing any individual record to be proven a member of the set with a short proof rather than a full re-scan. Together these primitives let a verifier detect modification cheaply and locally (Lu, 2019b).

Layered on top of these primitives is a mechanism for agreement. In a distributed setting, multiple parties must concur on the order and validity of new records without a central arbiter. Consensus protocols provide this agreement, and they vary widely: proof-based protocols tie the right to append to demonstrable expenditure of a resource, while voting-based protocols reach agreement through coordinated message exchange among known validators. The choice of protocol governs much of a system's resilience and throughput profile, and a substantial body of work now compares these protocols along security, scalability, and fault-tolerance lines (Xiao et al., 2020). Consensus is what turns a tamper-evident data structure into a tamper-evident system, because it ensures that the verifiable history is also a shared history.

It is useful to separate three properties that are frequently merged. Tamper-evidence is the ability to detect unauthorized change. Provenance is the ability to reconstruct the origin and transformation history of a data item. Resilience is the ability of the system to keep operating and to preserve its guarantees despite faults or attacks. A design can be strong on one and weak on another. Native provenance, for instance, is not automatic: querying the full history of a value in many ledgers requires replaying all transactions, which is impractical online, and dedicated provenance layers have been proposed to expose lineage efficiently to applications at runtime (Ruan et al., 2019; Ramachandran & Kantarcioglu, 2018).

Provenance becomes especially demanding when the data originates outside the ledger, as it does in most cross-sector settings. Sensor readings, clinical observations, and customs declarations are generated by heterogeneous devices and organizations before they are ever committed, and a tamper-evident store can only attest to what it received, not to the truthfulness of the upstream source. Extensible provenance frameworks therefore treat the boundary between the physical or organizational source and the ledger as a first-class design concern, attaching verifiable metadata about origin and custody as data crosses that boundary (Sigwart et al., 2020; Lu & Xu, 2019).

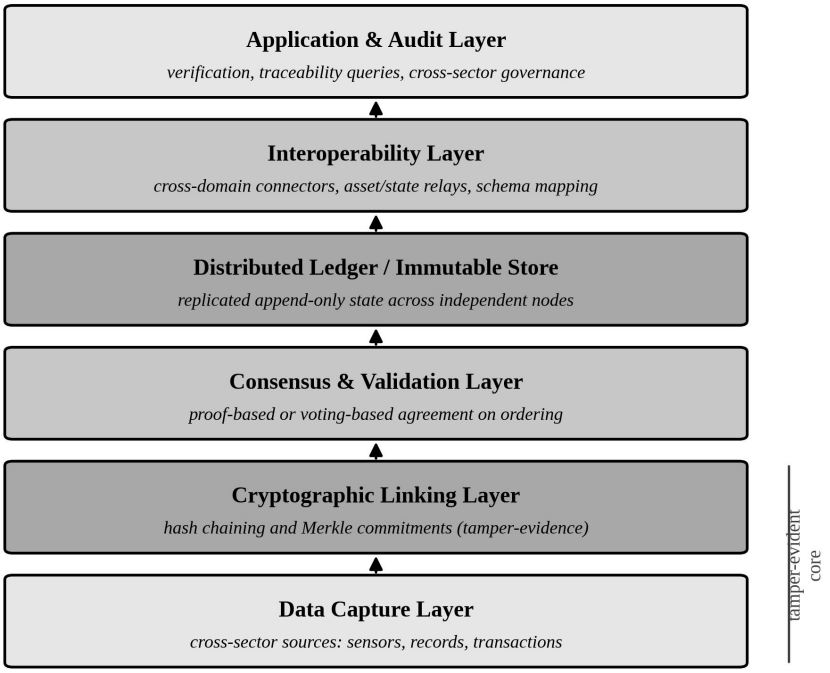

Figure 1. *Layered reference model of a tamper-evident data architecture, with cryptographic linking and replicated storage forming the tamper-evident core.*

Figure 1 presents the layered architecture that organizes the rest of our analysis. At the base, a data-capture layer ingests records from cross-sector sources. Above it, the cryptographic-linking layer applies hash chaining and Merkle commitments to render the captured data tamper-evident. The consensus-and-validation layer establishes agreement on ordering and validity, and the distributed-ledger layer maintains the replicated, append-only state across independent nodes. An interoperability layer connects the ledger to other domains and systems, and an application-and-audit layer exposes verification, traceability queries, and governance functions to end users. The figure deliberately places cryptographic linking and replicated storage at the centre, because these layers constitute the tamper-evident core on which every higher-level trust claim depends.

Two cautions follow from this layered view. First, tamper-evidence at the storage layer says nothing about correctness at the application layer: a smart contract that faithfully executes flawed logic will produce tamper-evident but incorrect outcomes, and the security of programmable layers has its own substantial failure surface (Atzei et al., 2017). Second, transparency and confidentiality pull in opposite directions. A ledger that is fully transparent maximizes auditability but can expose sensitive information, while techniques that restore confidentiality—encryption, selective disclosure, and zero-knowledge proofs—add complexity and cost. Privacy is therefore not a feature that can be assumed; it is a property that must be engineered against the grain of transparency (Feng et al., 2019).

Resilience in these systems is bounded by classical results in distributed computing. Byzantine fault tolerance—the ability to maintain agreement despite components that fail arbitrarily or behave maliciously—comes at a cost in communication and, depending on the protocol, in scalability, and the achievable fault threshold shapes how much adversarial behaviour a network can absorb before its guarantees degrade (Xiao et al., 2020). A related tension, familiar from distributed databases, is that a system cannot simultaneously maximize consistency, availability, and tolerance to network partitions; designers must choose which to relax under stress. Tamper-evident architectures

inherit these constraints directly: a widely replicated public ledger favours resilience and consistency at the expense of latency, whereas a centralized log favours availability and throughput while concentrating the very control that decentralization was meant to remove.

3. A Comparative Framework

The comparative framework distinguishes six architectural patterns that recur across deployments. The public permissionless ledger admits any participant and secures itself through open consensus; it maximizes openness and censorship resistance but pays in throughput, latency, and energy or capital cost. The permissioned or consortium ledger restricts validation to known, partially trusted organizations, improving performance and governance while reducing the degree of decentralization. The hybrid public–private design keeps sensitive operations in a permissioned domain while selectively anchoring or disclosing information to a broader network, trading additional complexity for flexibility. The directed-acyclic-graph ledger replaces the linear chain with a graph of confirmations to raise concurrency and throughput, often at the cost of maturity and interoperability.

The remaining two patterns are frequently overlooked in trust discussions but are important in practice. Layer-two and sidechain-anchored systems move most activity to a faster auxiliary system and periodically commit compressed evidence to a more secure base ledger, improving throughput and cross-domain reach while introducing new bridging assumptions. The centralized tamper-evident log retains a single operator but publishes periodic cryptographic commitments—Merkle roots or hash-chain checkpoints—so that the operator cannot silently rewrite history; it offers the efficiency, privacy, and governance clarity of a conventional database while still providing detection of tampering. Treating this last pattern as a legitimate member of the design space is important, because for many regulated applications it is the most pragmatic option.

Table 1. *Tamper-evident architectural patterns compared in this study.*

Pattern	Core orientation	Primary strength	Primary limitation	Representative cross-sector use
Public permissionless	Open, trustless replication	Resilience; censorship resistance	Throughput; privacy; cost	Open finance; public notarization
Permissioned / consortium	Governed multi-party ledger	Balanced governance and privacy	Reduced decentralization	Inter-firm data sharing; trade
Hybrid (public-private)	Selective disclosure	Flexibility; tunable openness	Design and integration complexity	Energy data; regulated reporting
DAG-based	Concurrent confirmation	High throughput	Maturity; interoperability	High-volume IoT telemetry
Layer-2 / sidechain-anchored	Off-base execution, anchored evidence	Throughput; cross-domain reach	Bridging trust assumptions	Micropayments; scaling settlement
Centralized tamper-evident log	Single operator, published commitments	Efficiency; privacy; governance	Resilience; single custodian	Internal audit; clinical records

We evaluate these patterns along seven dimensions chosen to capture the properties that matter for cross-sector trust. Table 1 summarizes the patterns and their characteristic orientation, while Table 2 defines the dimensions.

Traceability granularity measures how precisely the origin and transformation of a data item can be reconstructed. Interoperability measures how readily the system exchanges data and value with other domains, a property that has become central as organizations build heterogeneous, multi-ledger landscapes (Belchior et al., 2021). System resilience measures fault and attack tolerance and the preservation of guarantees under stress. Throughput captures sustainable transaction volume and latency, the dimension most often cited as the binding constraint on adoption (Zhou et al., 2020; Xie et al., 2019).

The remaining three dimensions capture properties that are easy to neglect until they fail. Privacy measures the system's ability to protect sensitive content while preserving verifiability, balancing the transparency that enables audit against the confidentiality that regulation and competition demand (Feng et al., 2019). Governance clarity measures how well-defined the rules are for participation, upgrade, dispute resolution, and the correction of erroneous entries; weak governance is a frequent cause of failed or contested deployments. Cost-efficiency measures the resource, operational, and integration burden of running the system at the required scale. Identity and access control cut across several dimensions, since the credibility of any traceability or governance claim ultimately depends on knowing which real-world entity stands behind a key (Mühle et al., 2018).

Table 2. *Evaluation dimensions used in the comparative coding.*

Dimension	Meaning in this study	A high score implies
Traceability granularity	Precision of origin/transformation reconstruction	Fine-grained provenance and audit
Interoperability	Ease of cross-domain data and value exchange	Strong fit for multi-ledger landscapes
System resilience	Fault and attack tolerance under stress	Robust continuity of guarantees
Throughput	Sustainable transaction volume and latency	Suitability for high-volume settlement
Privacy	Protection of sensitive content with verifiability	Fit for regulated, confidential data
Governance clarity	Defined rules for participation, upgrade, dispute	Lower risk of contested operation
Cost-efficiency	Resource, operational and integration burden	Lower total cost at required scale

The scoring is deliberately comparative rather than absolute. A score of five on a dimension does not denote perfection in any abstract sense; it denotes strong relative suitability within this set of six patterns. This is a structured qualitative coding device, not a claim of measured performance for any particular product, and it is anchored in the documented behaviour of representative implementations rather than in a single benchmark. The approach mirrors established practice in methodological benchmarking, where the contribution lies in a transparent and reproducible comparison framework rather than in an assertion of immutable numerical truth. Readers who weight the dimensions differently can recompute the scenario rankings using the same matrix, which is precisely the point of making the coding explicit.

Framing traceability, interoperability, and resilience as explicit dimensions also clarifies a conceptual claim that motivates the whole framework. These three are too often assumed to follow automatically from adopting a ledger. In reality they are in tension: maximizing resilience through wide, open replication tends to depress throughput and complicate privacy; maximizing interoperability through bridges introduces new trust assumptions and attack surface; and maximizing fine-grained traceability multiplies the data that must be stored, replicated, and protected. A

framework that names these dimensions and forces an explicit weighting is therefore not a bureaucratic overlay; it is the mechanism by which a designer confronts trade-offs that would otherwise be discovered late and at high cost (Leng et al., 2022; Chen et al., 2024).

4. Cross-Sector Analysis and Results

Figure 2 presents the coded suitability matrix for the six patterns across the seven dimensions. Several contrasts are immediately visible. The public permissionless ledger scores highest on resilience and strongly on traceability and auditability, but lowest on throughput, privacy, and governance clarity, reflecting the price of open participation. The centralized tamper-evident log occupies almost the opposite corner: high on throughput, privacy, governance, and cost-efficiency, but lowest on resilience and only moderate on interoperability. The permissioned and hybrid patterns sit between these poles and exhibit the most balanced profiles, which is one reason they dominate regulated enterprise deployments. The directed-acyclic-graph and layer-two patterns are throughput-oriented, with the latter scoring better on interoperability because anchoring naturally spans systems.

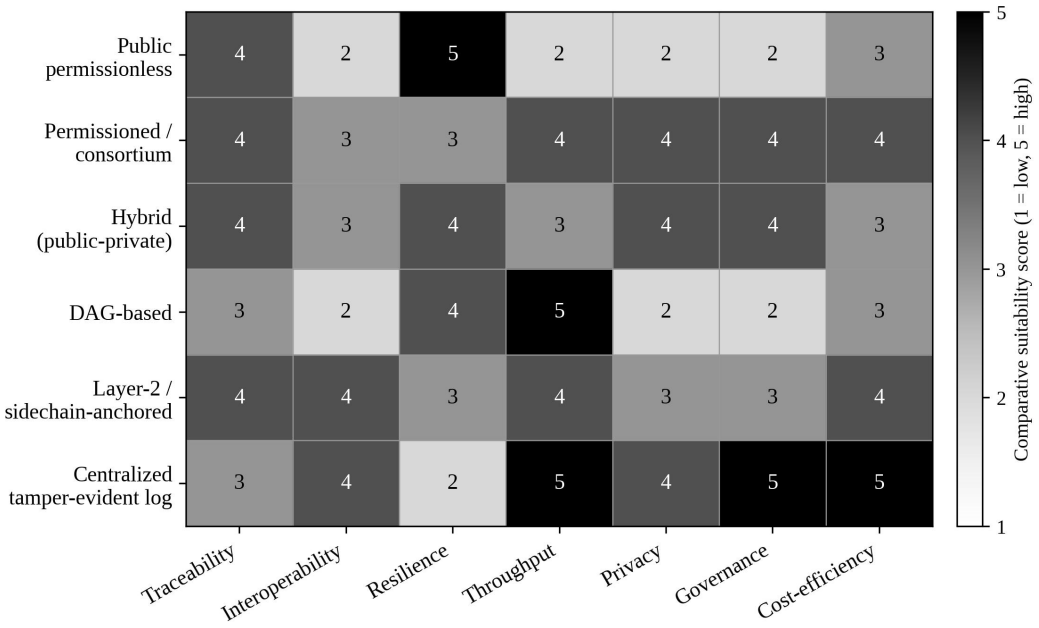

Figure 2. *Comparative coding of six tamper-evident architectures across seven trust dimensions (darker cells denote higher relative suitability).*

Reading the matrix by column is equally instructive. No single column is dominated by one pattern across all others, which is the central empirical message of the analysis: there is no universally best tamper-evident architecture. Resilience favours open replication, privacy and governance favour controlled membership, and throughput favours either centralization or graph-based concurrency. Because these columns point in different directions, any claim that one architecture is simply superior conceals an implicit and often unexamined weighting of the dimensions. The matrix makes that weighting visible and therefore debatable.

To translate the matrix into decision-relevant guidance, we weight the dimensions for three representative cross-sector scenarios and compute a weighted suitability score for each pattern. Table 3 records the weighting logic. In supply-chain traceability, the highest weights fall on traceability granularity, interoperability, and cost-efficiency, because the dominant requirement is to follow goods across many organizations affordably. In healthcare data exchange, weight shifts to privacy, governance clarity, resilience, and traceability, reflecting the sensitivity of clinical data and the regulatory environment in which it moves (Agbo et al., 2019). In energy peer-to-peer trading, weight

concentrates on throughput, resilience, interoperability, and cost-efficiency, because many small participants settle frequent transactions in near real time.

Table 3. *Scenario weighting logic and leading patterns by weighted suitability.*

Scenario	Most heavily weighted dimensions	Leading pattern	Runner-up
Supply-chain traceability	Traceability; interoperability; cost-efficiency	Centralized tamper-evident log	Permissioned / consortium
Healthcare data exchange	Privacy; governance; resilience; traceability	Permissioned / consortium	Hybrid (public-private)
Energy peer-to-peer trading	Throughput; resilience; interoperability; cost	Centralized tamper-evident log	Permissioned / consortium

Figure 3 reports the resulting scenario-weighted scores. In the supply-chain scenario, the centralized tamper-evident log and the permissioned ledger lead, with hybrid and layer-two designs close behind; the public permissionless ledger trails because its throughput and cost penalties outweigh its resilience advantage when goods tracking, not censorship resistance, is the objective. In the healthcare scenario, the permissioned and hybrid patterns lead clearly, since their governance and privacy strengths align with the dimensions that receive the most weight, and the centralized log remains competitive where a single accountable custodian is acceptable. In the energy scenario, the centralized log and permissioned ledger again score well, but the directed-acyclic-graph and layer-two patterns rise relative to other scenarios because their throughput orientation matches the settlement workload.

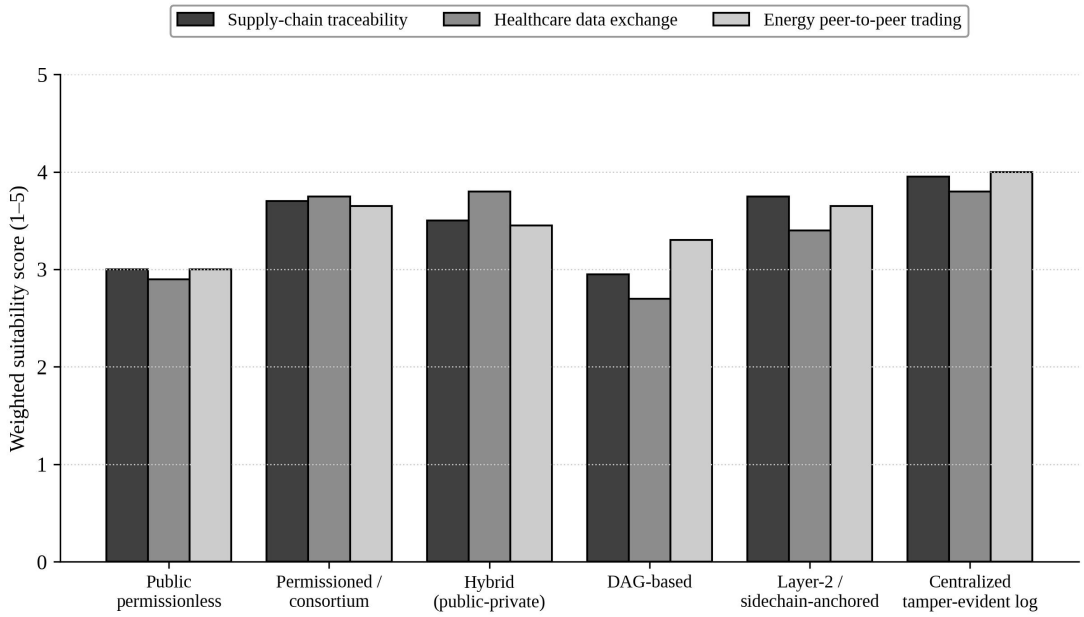

Figure 3. *Scenario-weighted suitability scores for the six tamper-evident architectures across three cross-sector settings.*

These results should be read as illustrative of a method rather than as verdicts on specific products, but they capture patterns that practitioners will recognize. The recurring strength of the permissioned and centralized patterns across scenarios reflects a real tendency in deployed systems: when a small set of accountable organizations already exists, the marginal benefit of open, permissionless replication is often outweighed by its costs, and tamper-evidence can be

obtained more cheaply. Conversely, the public ledger's advantage is greatest precisely where no such trusted set exists and where resistance to unilateral control is itself the trust requirement, a configuration more typical of open financial settings than of intra-industry data sharing (Xu et al., 2024).

Sector context further refines these rankings. In food and pharmaceutical supply chains, traceability granularity interacts with provenance at the physical boundary, because a tamper-evident record of a shipment is only as trustworthy as the device or party that reported it; this is why provenance frameworks that bind origin metadata at the point of capture are decisive complements to the ledger itself (Sigwart et al., 2020; Ramachandran & Kantarcioglu, 2018). In healthcare, interoperability across institutional silos and rigorous identity binding are persistent obstacles, and self-sovereign identity models have been proposed to give individuals controllable, verifiable credentials across providers (Mühle et al., 2018; Agbo et al., 2019). In energy and industrial settings, the integration of ledgers with large fleets of constrained devices raises both scalability and security concerns that middleware-oriented architectures attempt to manage (Leng et al., 2022; Xu et al., 2021).

Two further sectors illustrate the endpoints of the design space. In finance, and especially in decentralized finance, value itself is recorded and moved on programmable, often permissionless infrastructure, so the trust requirement is dominated by resilience, openness, and resistance to unilateral control; here the public permissionless pattern earns its throughput and privacy penalties because no single custodian is acceptable, and the security of the programmable layer becomes systemically important (Xu et al., 2024; Atzei et al., 2017). In public administration, by contrast, records such as land titles, licences, and identity attributes demand strong governance, accountability, and the ability to correct errors under legal authority, which favours permissioned or hybrid designs, or centralized tamper-evident logs operated by an accountable authority. The contrast underscores the paper's central claim: the same tamper-evident guarantee is delivered by very different architectures, and the appropriate one is fixed only once the dominant trust property is named.

A further analytical observation concerns interoperability as a system-level rather than pairwise property. As soon as a sector operates more than one ledger—public for settlement, permissioned for sensitive operations, centralized logs for internal audit—the trust of the whole depends on the connectors between them. Bridges and relays inherit the weakest security of the systems they join and add their own, so interoperability gains can quietly erode resilience. The framework captures this by scoring interoperability separately, but the deeper lesson is that cross-sector digital trust is increasingly a property of architectures of architectures, and that the connective tissue deserves the same scrutiny as the ledgers themselves (Belchior et al., 2021).

5. Discussion

The analysis carries a clear methodological implication: tamper-evident architecture should be chosen before the application is built, and the choice should be driven by an explicit weighting of trust dimensions rather than by platform familiarity. The recurring tension among resilience, throughput, privacy, and interoperability is a structural feature, not a temporary limitation that the next platform will dissolve. Naming and weighting the dimensions converts an implicit gamble into a defensible design decision, and it makes the decision reproducible, since a reviewer can recompute the ranking under different weights. This reframing is the principal contribution of the paper, and it generalizes beyond the six patterns examined here.

A second implication concerns governance and the management of error. Because tamper-evidence makes silent change detectable rather than impossible, every serious deployment needs an explicit, auditable process for correcting validly recorded but erroneous data, for rotating compromised keys, and for evolving the protocol. Where such governance is weak, the very property that creates trust—detectable permanence—becomes a liability, as incorrect entries cannot be quietly fixed and disputes escalate. Analytical and decision-support tooling can help operators

monitor these processes and reason about concentration of control and systemic risk, linking architectural choices to measurable governance indicators (Lu et al., 2024; Zhang & Lu, 2021).

Tamper-evident architectures also do not exist in isolation; they are converging with adjacent technologies that reshape their trade-offs. The integration of ledgers with the Internet of Things extends tamper-evidence to the data-capture boundary but stresses scalability and device security (Xu et al., 2021; Lu & Xu, 2019). Convergence with artificial intelligence introduces both an opportunity, in the form of automated anomaly detection over verifiable data, and a risk, in the form of models whose training data and behaviour themselves require provenance and audit (Lu, 2019a; Chen et al., 2024). In industrial and cyber-physical settings, tamper-evident data is becoming a foundation for trustworthy automation rather than an end in itself (Lu, 2017a; Lu, 2017b).

The communications substrate matters as well. Edge and next-generation network research anticipates dense, low-latency environments in which cross-sector data is generated and settled close to its source, which favours architectures that can push verification toward the edge and reconcile it centrally (Lu & Zheng, 2020; Lu & Ning, 2020). Resource-allocation and market mechanisms developed for cloud and service settings offer transferable insight for the auction and settlement logic of decentralized energy and data marketplaces (Lu et al., 2020). And in finance, the migration of value onto programmable, permissionless infrastructure has made the governance and security of tamper-evident systems a matter of systemic importance rather than a niche concern (Xu et al., 2024).

A forward-looking caveat concerns the durability of the cryptographic core. Tamper-evidence ultimately rests on the hardness assumptions behind hash functions and digital signatures, and advances in quantum computing threaten some of these assumptions, which has motivated research into quantum-resistant primitives and into the broader implications of quantum capability for information infrastructure (Lu et al., 2023; Lu & Yang, 2024). Because the records written today may need to remain verifiable for decades, architectural planning should treat cryptographic agility—the ability to migrate to stronger primitives without rewriting history—as a design dimension in its own right, even though it lies beyond the seven dimensions scored here.

A further implication concerns standardization. As cross-sector deployments multiply, the absence of shared standards for data models, identifiers, and cross-ledger messaging becomes a primary obstacle to interoperability, forcing organizations to build bespoke and fragile connectors between systems that were never designed to interoperate (Belchior et al., 2021). Standardized interfaces, verifiable-credential formats, and common provenance schemas would lower integration cost and reduce the attack surface introduced by ad hoc bridges, while also making cross-sector audit more tractable. This is as much a governance and policy question as a technical one, because standards must be agreed among stakeholders with divergent incentives, and their adoption depends on regulatory encouragement as well as engineering merit (Chen et al., 2024). Treating standardization as part of the architecture, rather than as an afterthought, is therefore essential to realizing cross-sector digital trust at scale.

Several limitations bound these conclusions. The comparative coding compresses rich architectural behaviour into ordinal scores, and reasonable experts may assign different values; the framework accommodates this by exposing the weights, but it does not eliminate judgement. The six patterns are archetypes, and real systems often blend them, so a given deployment may inherit a mixture of the profiles described here. The three scenarios, though representative, do not exhaust the space of cross-sector applications, and sectors such as public administration, education, and digital identity would warrant their own weightings. Finally, the analysis evaluates architectures as designs rather than measuring fielded implementations; a natural extension would link the coded dimensions to empirical performance and security outcomes across replicated case studies, complementing scalability and management-analytic evidence already emerging in the literature (Zhou et al., 2020; Lu et al., 2024).

6. Conclusion

This article has argued that digital trust is increasingly a property of data architectures and that tamper-evidence—the detectability of unauthorized change—is the property that lets mutually distrustful parties share a credible record. Treating distributed-ledger technology as a single artefact obscures a wide design space, and we have instead distinguished six recurring patterns and assessed them through a transparent comparative coding across seven dimensions: traceability, interoperability, resilience, throughput, privacy, governance, and cost-efficiency.

The central finding is that no architecture dominates. Public permissionless ledgers maximize resilience and auditability but are constrained on throughput, privacy, and governance; centralized tamper-evident logs invert that profile; and permissioned and hybrid designs offer the most balanced fit for regulated, cross-organizational data exchange. Scenario-weighted analysis for supply-chain traceability, healthcare data exchange, and energy peer-to-peer trading shows that the best choice shifts with the trust property a sector prioritizes, which is exactly why the dimensions must be weighted explicitly rather than assumed.

The broader message for researchers, developers, and policymakers is that traceability, interoperability, and resilience should be elevated from assumed by-products to explicit, weighted design objectives, and that the connective tissue between heterogeneous ledgers deserves the same scrutiny as the ledgers themselves. Future work should extend the framework to additional sectors, link its coded dimensions to empirical outcomes, and incorporate cryptographic agility so that architectures designed today remain trustworthy as both threats and technologies evolve.

Declaration of AI-assisted language editing

During the preparation of this manuscript, language-model assistance was used only for English-language editing and document organisation. The authors designed the framework, performed the comparative coding and analysis, and reviewed and take full responsibility for the content, tables, figures, and interpretations.

References

Agbo, C. C., Mahmoud, Q. H., & Eklund, J. M. (2019). Blockchain technology in healthcare: A systematic review. Healthcare, 7(2), 56. https://doi.org/10.3390/healthcare7020056

Atzei, N., Bartoletti, M., & Cimoli, T. (2017). A survey of attacks on Ethereum smart contracts (SoK). In Principles of Security and Trust (POST 2017), LNCS 10204 (pp. 164–186). Springer. https://doi.org/10.1007/978-3-662-54455-6_8

Belchior, R., Vasconcelos, A., Guerreiro, S., & Correia, M. (2021). A survey on blockchain interoperability: Past, present, and future trends. ACM Computing Surveys, 54(8), Article 168. https://doi.org/10.1145/3471140

Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. Information Systems Frontiers, 26(5), 1715–1729. https://doi.org/10.1007/s10796-022-10248-7

Feng, Q., He, D., Zeadally, S., Khan, M. K., & Kumar, N. (2019). A survey on privacy protection in blockchain system. Journal of Network and Computer Applications, 126, 45–58. https://doi.org/10.1016/j.jnca.2018.10.020

Leng, J., Chen, Z., Huang, Z., Zhu, X., Su, H., Lin, Z., & Zhang, D. (2022). Secure blockchain middleware for decentralized IIoT towards Industry 5.0: A review of architecture, enablers, challenges, and directions. Machines, 10(10), 858. https://doi.org/10.3390/machines10100858

Lu, Y. (2017a). Cyber physical system (CPS)-based Industry 4.0: A survey. Journal of Industrial Integration and Management, 2(3), 1750014. https://doi.org/10.1142/S2424862217500142

Lu, Y. (2017b). Industry 4.0: A survey on technologies, applications and open research issues. Journal of Industrial Information Integration, 6, 1–10. https://doi.org/10.1016/j.jii.2017.04.005

Lu, Y. (2018). Blockchain and the related issues: A review of current research topics. Journal of Management Analytics, 5(4), 231–255. https://doi.org/10.1080/23270012.2018.1516523

Lu, Y. (2019a). Artificial intelligence: A survey on evolution, models, applications and future trends. Journal of Management Analytics, 6(1), 1–29. https://doi.org/10.1080/23270012.2019.1570365

Lu, Y. (2019b). The blockchain: State-of-the-art and research challenges. Journal of Industrial Information Integration, 15, 80–90. https://doi.org/10.1016/j.jii.2019.04.002

Lu, Y. (2022). Implementing blockchain in information systems: A review. Enterprise Information Systems, 16(12), 1876–1907. https://doi.org/10.1080/17517575.2021.2008513

Lu, Y., & Ning, X. (2020). A vision of 6G–5G's successor. Journal of Management Analytics, 7(3), 301–320. https://doi.org/10.1080/23270012.2020.1802622

Lu, Y., & Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. IEEE Internet of Things Journal, 6(2), 2103–2115. https://doi.org/10.1109/JIOT.2018.2869847

Lu, Y., & Yang, J. (2024). Quantum financing system: A survey on quantum algorithms, potential scenarios and open research issues. Journal of Industrial Information Integration, 41, 100663. https://doi.org/10.1016/j.jii.2024.100663

Lu, Y., & Zheng, X. (2020). 6G: A survey on technologies, scenarios, challenges, and the related issues. Journal of Industrial Information Integration, 19, 100158. https://doi.org/10.1016/j.jii.2020.100158

Lu, Y., Pisarenko, Z. V., Yang, L., & Ye, C. (2024). Advancing decision-making: The role of management analytics in modern business practices. Nanotechnologies in Construction, 16(5), 431–440. https://doi.org/10.15828/2075-8545-2024-16-5-431-440

Lu, Y., Sigov, A. S., Ratkin, L., Ivanov, L. A., & Zuo, M. (2023). Quantum computing and industrial information integration: A review. Journal of Industrial Information Integration, 35, 100511. https://doi.org/10.1016/j.jii.2023.100511

Lu, Y., Zheng, X., Li, L., & Xu, L. D. (2020). Pricing the cloud: A QoS-based auction approach. Enterprise Information Systems, 14(3), 334–351. https://doi.org/10.1080/17517575.2019.1669827

Mühle, A., Grüner, A., Gayvoronskaya, T., & Meinel, C. (2018). A survey on essential components of a self-sovereign identity. Computer Science Review, 30, 80–86. https://doi.org/10.1016/j.cosrev.2018.10.002

Ramachandran, A., & Kantarcioglu, M. (2018). SmartProvenance: A distributed, blockchain based data provenance system. In Proceedings of the 8th ACM Conference on Data and Application Security and Privacy (CODASPY '18) (pp. 35–42). https://doi.org/10.1145/3176258.3176333

Ruan, P., Chen, G., Dinh, T. T. A., Lin, Q., Ooi, B. C., & Zhang, M. (2019). Fine-grained, secure and efficient data provenance on blockchain systems. Proceedings of the VLDB Endowment, 12(9), 975–988. https://doi.org/10.14778/3329772.3329775

Sigwart, M., Borkowski, M., Peise, M., Schulte, S., & Tai, S. (2020). A secure and extensible blockchain-based data provenance framework for the Internet of Things. Personal and Ubiquitous Computing. https://doi.org/10.1007/s00779-020-01417-z

Xiao, Y., Zhang, N., Lou, W., & Hou, Y. T. (2020). A survey of distributed consensus protocols for blockchain networks. IEEE Communications Surveys & Tutorials, 22(2), 1432–1465. https://doi.org/10.1109/COMST.2020.2969706

Xie, J., Yu, F. R., Huang, T., Xie, R., Liu, J., & Liu, Y. (2019). A survey on the scalability of blockchain systems. IEEE Network, 33(5), 166–173. https://doi.org/10.1109/MNET.001.1800290

Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. IEEE Internet of Things Journal, 8(13), 10452–10473. https://doi.org/10.1109/JIOT.2021.3060508

Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. Enterprise Information Systems, 18(9). https://doi.org/10.1080/17517575.2024.2397630

Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. Journal of Industrial Information Integration, 23, 100224. https://doi.org/10.1016/j.jii.2021.100224

Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. Enterprise Information Systems, 16(12), 1939895. https://doi.org/10.1080/17517575.2021.1939895

Zhou, Q., Huang, H., Zheng, Z., & Bian, J. (2020). Solutions to scalability of blockchain: A survey. IEEE Access, 8, 16440–16455. https://doi.org/10.1109/ACCESS.2020.2967218