

Blockchain-Based Supply Chain Transparency for Agricultural Produce: A Systematic Literature Review and System Design

Dr. V. Janani¹, Chandana JA², Shravani PV², & Vaishnavi S²

¹Assistant Professor and HOD, Department of Information Science and Engineering
Sri Sairam College of Engineering, Bengaluru

²Final Year BE (Information Science and Engineering), Sri Sairam College of Engineering, Bengaluru

DOI: doi.org/10.34293/iejcsa.v4i5.173

Abstract - Agricultural supply chains involve many independent participants—farmers, aggregators, cold-storage operators, distributors, and retailers—each of whom typically records custody, quality, and pricing information in a separate, non-interoperable system, if at all. This fragmentation makes it difficult for any single participant, and nearly impossible for the end consumer, to verify where a given batch of produce originated, how it was handled in transit, or whether the claims printed on its label (organic, fair-trade, pesticide-free) are actually true. This paper surveys eight recent works on blockchain-based agricultural traceability, spanning farm-to-fork provenance frameworks, smart-contract-driven payment and quality-verification systems, consensus-mechanism surveys, IoT-integrated cold-chain monitoring, decentralized organic certification, hybrid blockchain-cloud architectures, consumer-facing QR/RFID applications, and cross-border interoperability studies. Read together, these works show a consistent movement away from siloed, paper-or database-based record-keeping and toward tamper-evident, cryptographically verifiable custody chains. At the same time, nearly every reviewed system still depends on an unverified, manual data-entry step at the point where a physical batch first enters the digital ledger—the so-called last-mile trust gap—and few works give ordinary consumers or smallholder farmers an interface simple enough to use without technical training. Building on these findings, the paper presents the design of a Blockchain-Based Agricultural Supply-Chain Transparency Platform: a three-tier architecture combining an IoT and RFID/QR data-capture layer at each custody transfer, a permissioned Hyperledger Fabric ledger that records hashed provenance and quality events with off-chain bulk storage on IPFS, and a consumer- and farmer-facing application layer that renders the resulting custody chain as a scannable, human-readable provenance trail. The architecture is intended to move agricultural supply chains from opaque, trust-based coordination toward verifiable, evidence-based transparency.

Keywords: Blockchain, Agricultural Supply Chain, Food Traceability, Smart Contracts, Hyperledger Fabric, IoT, Provenance Verification, Food Fraud Prevention

INTRODUCTION

The journey of a unit of agricultural produce from a field to a consumer's shopping bag typically passes through five or more independent custodians, each of whom records what happened under their watch in a system that the next custodian in the chain cannot see, let alone verify. A farmer may log a harvest date in a notebook or a village cooperative's spreadsheet; an aggregator may re-weigh and repackage the same batch under a new lot number with no durable link back to the original; a distributor's warehouse management

system tracks temperature and time in storage but does not expose that data to the retailer, let alone the consumer. By the time a package reaches a shelf, the label's claims about origin, handling, or certification status rest almost entirely on the word of whichever party printed it, because no participant earlier in the chain, and certainly no consumer at the end of it, has any practical way to independently check.

This absence of end-to-end verifiability is not merely an inconvenience; it has direct economic and public-health consequences. Certification fraud, unauthorized substitution of a cheaper commodity for a premium one, and the deliberate mislabeling of country of origin all persist in large part because the systems that would need to detect them are fragmented across mutually distrusting parties who have no shared, tamper-evident record to check against. Foodborne-illness outbreaks compound the same structural problem: when contamination is discovered downstream, tracing it back to a specific farm or lot can take investigators days or weeks under a paper- and spreadsheet-based recall process, during which the contaminated product often remains on shelves. The economic stakes are considerable: global estimates of food fraud losses run into the tens of billions of dollars annually, and premium-certified categories—organic, fair-trade, single-origin—carry price premiums specifically because the certification is trusted, which means the entire premium is at risk whenever that trust cannot be independently verified. A supply chain that cannot prove provenance also cannot isolate a contamination event quickly, which converts what might have been a narrow, single-lot recall into a broad, category-wide one that damages far more stock, and far more trust, than the original defect warranted.

A parallel and increasingly practical development is the maturation of distributed-ledger technology into a form that no longer requires every participant to trust a single central database operator. Permissioned blockchain platforms allow a defined consortium of supply-chain participants—farmers, cooperatives, logistics providers, certifiers, and retailers—to write to a shared ledger that none of them individually controls, while smart contracts encode the business rules (a quality threshold, a payment release condition, a certification requirement) that previously depended on manual verification or bilateral trust. Blockchain-protocol work tends to focus on what happens once an event has already been digitized, while IoT and sensing work focuses on how a physical event (a temperature excursion, a location change, a weight reading) is captured in the first place. Few of the works reviewed here attempt to connect the two rigorously, even though the integrity of everything downstream of ingestion depends entirely on how faithfully that first digitization step reflects physical reality.

This paper pursues four objectives. First, to organize the reviewed approaches by their primary design emphasis—provenance tracking, smart-contract automation, consensus and platform selection, IoT sensor integration, certification, hybrid architecture, consumer interface, or cross-chain interoperability. Second, to compare the advantages and structural limitations reported for each approach. Third, to identify the gaps that most

directly constrain a deployable, farm-to-fork transparency platform. Fourth, to translate those gaps into concrete architectural decisions for the proposed platform. The remainder of the paper is organized as follows: the Literature Review section outlines the technical foundations, explains how the eight sources were selected and evaluated, works through them by theme, and compares them; the Research Gap and Problem Statement sections distill the resulting gaps into a precise problem; and the Proposed System, System Architecture, and Methodology/Workflow sections use those gaps to justify and describe the proposed platform, followed by the Conclusion.

LITERATURE REVIEW

A. Background and Theoretical Foundations

Technical Evolution of Blockchain-Based Traceability

The approaches reviewed here fall loosely into three generations. The earliest generation digitizes existing paper or spreadsheet-based traceability records onto a private database or an early permissioned ledger, essentially replacing a filing cabinet with a shared table but leaving verification of what is entered entirely to the entering party. A second generation introduces smart contracts that encode business logic—a quality threshold that must be met before payment releases, a certification rule that must be satisfied before a batch can be labeled organic—directly into the ledger, so that some verification steps that previously required a human auditor happen automatically and consistently. The most recent generation, which this review emphasizes, layers IoT sensing, RFID/QR consumer interfaces, and cross-chain interoperability protocols on top of that automated core, so that physical events are captured with minimal manual entry and the resulting provenance trail is queryable by any participant, including the end consumer, regardless of which specific platform recorded a given event.

B. Review Methodology

The review follows a structured, two-step methodology designed to ensure that each source is assessed consistently and that the findings translate directly into design decisions.

Corpus and Search Scope

Eight studies addressing blockchain-based agricultural traceability, smart-contract payment and quality verification, consensus-mechanism selection for agri-blockchains, IoT-coldchain integration, decentralized certification, hybrid blockchain-cloud architectures, consumer-facing scanning applications, and cross-chain interoperability were selected for detailed review. Each was assessed for its stated architecture, its reported advantages, and the limitations the original authors themselves acknowledged.

Thematic Synthesis

The eight sources reviewed fall into four thematic groups: provenance and traceability frameworks; smart-contract-driven payment and quality assurance; IoT-

blockchain integration for physical-condition monitoring; and consensus, scalability, and interoperability studies. Reading across these groups rather than treating each source in isolation makes it easier to see where capabilities overlap, where they are redundant, and where a real deployment would need to combine two or more of them to be useful to a farmer, a regulator, and a consumer at once.

C. Provenance and Traceability Frameworks

A farm-to-fork traceability framework described in [1] records each custody transfer—harvest, aggregation, transport, retail receipt—as a discrete transaction on a permissioned ledger, giving any participant with read access a complete, chronologically ordered history of a batch from a single lot identifier. Its principal strength is immutability: once a transaction is committed, no single participant, including the original entrant, can quietly alter it, which removes an entire class of after-the-fact record tampering that plagues paper- and spreadsheet-based systems. Its central weakness, however, is the last-mile problem: the framework offers no described mechanism for verifying that data entered at harvest actually corresponds to physical reality, so a farmer or aggregator who enters an inaccurate weight or origin produces a permanent, cryptographically “trustworthy-looking” record of something that never happened.

A decentralized organic-certification approach in [5] extends the same ledger-based custody model specifically to certification claims, recording a certifying body’s approval as a signed transaction that travels with the batch and can be checked by any downstream party without contacting the certifier directly. This considerably shortens the certification-verification chain and makes fraudulent certification claims immediately detectable by comparing a label’s claim against the ledger, but the approach depends entirely on the certifying bodies themselves adopting the platform and correctly gating their approvals, and it does not address batches that are split, mixed, or repackaged after certification, which is common in bulk commodity handling.

A survey in [3] catalogues consensus mechanisms—proof-of-authority, practical Byzantine fault tolerance, and proof-of-stake variants—evaluated specifically for suitability in agricultural consortium settings, where participants are known and permissioned rather than anonymous. Its contribution is foundational rather than applied: it maps out which consensus families best balance transaction throughput against the low-connectivity, intermittent-participation conditions typical of rural nodes, but it does not attempt to show how a selected consensus mechanism would be wired into a working multi-tier system, nor does it address how consumer-facing queries would be served without requiring every reader to run a full node.

D. Smart-Contract-Driven Payment and Quality-Assurance Systems

A payment-automation system in [2] ties a smart contract’s execution to sensor-verified quality thresholds and delivery confirmation, releasing payment to a farmer

automatically once produce is received and its condition is attested on-chain, rather than after a manual invoice-and-approval cycle that can take days or weeks. This meaningfully improves farmer cash flow and removes a recurring source of payment disputes, but the tradeoff is that the quality thresholds and attestation logic must be specified correctly and immutably in advance; a poorly specified contract either releases payment for substandard produce or withholds it from acceptable produce due to a sensor miscalibration, and correcting a deployed contract is far more involved than editing a spreadsheet formula.

A related cross-border trade-finance approach, surveyed alongside [2] in the broader literature, layers letter-of-credit-style conditional payment logic on top of the same custody ledger, so that an importer's payment is escrowed and released only once customs, quality, and delivery conditions are all independently attested on-chain. This reduces reliance on traditional trade-finance intermediaries and shortens settlement time considerably, though it introduces new dependence on the reliability and honesty of whichever party operates each attesting sensor or inspection node.

E. IoT-Blockchain Integration for Cold-Chain and Sensor Data

An IoT-integrated cold-chain monitoring system in [4] attaches temperature and humidity loggers to shipments of perishable produce and writes periodic sensor readings, along with any threshold excursions, directly to the ledger as they occur, closing much of the last-mile trust gap for at least the in-transit portion of the chain. This gives downstream buyers verifiable, tamper-evident proof of cold-chain integrity rather than a paper log that could in principle be filled in after the fact, and it can trigger an automatic alert or contract penalty the moment a threshold is crossed. Its cost, however, scales directly with the number of sensor-equipped shipments, and the system's authors acknowledge that battery life and cellular or satellite connectivity in transit remain practical constraints, particularly for shipments originating from areas with limited network coverage.

A consumer-facing application in [7] closes the loop at the opposite end of the chain, allowing a shopper to scan a QR code or RFID tag on a retail package and retrieve a plain-language summary of the batch's recorded custody history, rendered from the same underlying ledger without requiring the consumer to understand blockchain terminology at all. This is the only reviewed work that treats the end consumer, rather than a supply-chain professional, as the primary interface user, and it demonstrably increases consumer trust and willingness to pay a premium for verifiably certified produce in the pilot deployments the authors describe; its limitation is that the mapping between a physical tag and its correct on-chain record still depends on correct tag application at the packaging stage, which is itself a manual step outside the ledger's guarantees.

F. Consensus, Scalability, and Interoperability Studies

A hybrid blockchain-cloud architecture in [6] addresses a different axis of the problem: how a system built for a single consortium can interoperate with the many other,

mutually incompatible blockchain platforms already adopted by different regional cooperatives and retailers. It keeps high-frequency, high-volume sensor and transaction data on a conventional cloud database while anchoring cryptographic hashes of that data onto a public or semi-public chain, giving the durability and low per-transaction cost of centralized storage alongside the tamper-evidence of a distributed ledger. Its cost grows with the volume of anchored data, and because the bulk of the data still lives off-chain, the system's guarantees are only as strong as the cloud operator's own integrity and uptime, which reintroduces a degree of the centralized trust the approach was meant to reduce.

A dedicated interoperability study [8] catalogues the practical difficulties of reconciling different consortiums' choice of platform—Hyperledger Fabric in one region, a public Ethereum-compatible chain in another—when a single shipment crosses between them, and proposes cross-chain bridging and shared identifier standards as a partial remedy. This work is foundational rather than applied, and the authors note that no production-scale cross-chain agricultural deployment yet exists to validate the proposed standards against real, high-volume trade flows.

G. Comparative Analysis

Two clear axes separate the reviewed works. The first is verification depth: [2], [4], and [5] attempt genuine, sensor- or rule-based verification of what is recorded, while [1] and [6] largely provide durable storage and replication without independently checking the accuracy of what enters the ledger. The second is end-user reach: only [7] extends the system's benefit to the end consumer rather than stopping at the level of supply-chain professionals, and only [2] closes the loop into an automatic operational action (a payment release) rather than stopping at a passive record. No single reviewed system combines automated, sensor-verified ingestion, smart-contract-automated payment and certification, and a consumer-facing interface within one architecture—each capability so far has been demonstrated in isolation.

A third, less obvious axis is accessibility at the farm-gate tier. Most of the reviewed systems assume a supply-chain participant comfortable operating enterprise software, managing cryptographic credentials, or interpreting a blockchain explorer. None of them treat the smallholder farmer's low-connectivity, low-literacy operating conditions as a first-class design constraint on the capture side, even though [7] demonstrates that a genuinely simple interface is achievable on the consumer side of the same chain. This suggests that farm-gate accessibility, rather than raw cryptographic assurance, may be the more immediate barrier separating the current body of agri-blockchain research from real adoption by the smallholder farmers who represent the majority of nodes in most agricultural supply chains.

A fourth pattern is that verification strength and deployment cost tend to move together rather than in opposition. The systems that most directly close the last-mile trust

gap—[4] through continuous sensing and [2] through automated attestation—also carry the highest per-shipment hardware and connectivity costs, whereas the systems with the lowest marginal cost—[1] and [6]—are precisely the ones that depend most heavily on unverified manual entry or a centrally operated anchor point. No reviewed system manages to combine strong ingestion-level verification with both low per-shipment cost and a genuinely accessible farm-gate interface at the same time, which is precisely the combination a practical, smallholder-inclusive transparency platform would need to strike.

Taken as a whole, this comparison suggests that the field has developed strong, individually validated components—immutable custody recording, automated payment, cold-chain sensing, certification attestation, and consumer-facing lookup—without yet producing an architecture that assembles them behind a single, farm-gate-accessible ingestion layer. This is less a criticism of any individual reviewed work, each of which addresses its own stated problem well, than an observation about the corpus as a whole: the next meaningful advance in agricultural supply-chain transparency is more likely to come from integration and accessible ingestion design than from further improving ledger-level cryptography in isolation.

Table 1 Summary of Reviewed Studies

Ref	Primary Approach	Key Reported Outcome / Limitation
[1]	Permissioned-ledger farm-to-fork custody recording	Tamper-evident history; no verification of data entered at ingestion
[2]	Smart-contract payment automation tied to quality attestation	Faster, dispute-resistant payment; brittle to sensor miscalibration
[3]	Consensus-mechanism survey for agricultural consortia	Maps consensus families to rural conditions; foundational, not applied
[4]	IoT cold-chain temperature/humidity logging to ledger	Verifiable cold-chain integrity; cost scales with sensor-equipped shipments
[5]	Decentralized organic-certification attestation	Fast fraud detection; assumes certifier adoption, weak to batch splitting
[6]	Hybrid blockchain-cloud anchoring architecture	Low per-transaction cost; reintroduces centralized-operator trust dependence
[7]	Consumer-facing QR/Rfid provenance lookup	High consumer trust and willingness to pay; tag-to-record binding still manual
[8]	Cross-chain interoperability standards for agri-trade	Addresses platform fragmentation; unvalidated at production trade volumes

RESEARCH GAP

Five gaps recur across the reviewed literature, in rough order of how foundational they are. Taken individually, each gap looks like a manageable engineering limitation; taken together, they describe a field where ledger-level guarantees have matured faster than the ingestion, interface, and accessibility safeguards needed to put those guarantees in front of an actual smallholder farmer or ordinary consumer.

- **The last-mile ingestion-trust gap.** Nearly every reviewed system—most explicitly [1]—assumes that data entered at the point of first digitization is accurate, with no described mechanism for automatically cross-checking a manually entered harvest date, origin, or weight against an independent physical signal. In practice, farm-gate and aggregator-level entry is done under time pressure by participants who are rarely trained in data hygiene or motivated to flag their own errors, so a ledger's immutability guarantees only that an inaccurate record will persist unchanged rather than that the record was accurate to begin with.
- **Absence of low-cost, farm-gate-accessible sensing.** IoT integration in the reviewed literature ([4]) is demonstrated primarily for in-transit cold-chain monitoring, where shipment values justify sensor cost; none of the reviewed works describe an affordable sensing or verification mechanism specifically for the farm-gate harvest event itself, which is the single point in the chain where an inaccurate entry does the most downstream damage.
- **Limited handling of batch splitting and repackaging.** The certification and traceability models in [1] and [5] both assume a batch retains a stable identity as it moves through the chain, yet real agricultural handling routinely splits, mixes, and repackages lots at aggregation and processing points; none of the reviewed systems describe a mechanism for preserving verifiable provenance across such a split with the same rigor as for an unsplit batch.
- **Platform fragmentation and weak interoperability.** Different consortiums, regions, and retailers adopt different underlying blockchain platforms, and cross-chain interoperability work remains largely theoretical, meaning that a shipment that crosses between two independently digitized supply chains can lose its verifiable custody trail at the handoff point even though each individual chain is internally sound.
- **No unified interface spanning both ends of the chain.** Consumer-facing interfaces ([7]) and farmer- or administrator-facing systems ([1], [2], [6]) are demonstrated separately; none of the reviewed works describe a single platform that offers a genuinely simple capture interface for a smallholder farmer and a genuinely simple lookup interface for a consumer from the same underlying ledger, which forces any real deployment to stitch together components built by different teams under different assumptions about their users.

PROBLEM STATEMENT

Three recurring obstacles surface across the literature considered here. The first is data fragmentation across custodians: each participant in the chain typically maintains its own record system, with no shared schema or identifier linking a batch as it changes hands, so that provenance information that exists somewhere in the chain is effectively invisible to

everyone else in it. The second is the last-mile trust gap at the point of physical-to-digital conversion: even a perfectly tamper-proof ledger only records what it is told, so if a farmer or aggregator enters an inaccurate harvest date, origin, or weight at the point of first digitization, that inaccuracy becomes a permanently and convincingly recorded falsehood rather than a correctable spreadsheet error. The third is the accessibility gap between the technical sophistication of most blockchain tooling and the practical, often low-connectivity conditions under which smallholder farmers and rural aggregators actually operate.

These three obstacles compound rather than operate independently. A consortium that solves data fragmentation by adopting a shared ledger still inherits whatever inaccuracy enters at the farm gate, and a system that solves the trust gap through IoT sensing at every node still fails in practice if the sensing hardware and connectivity it requires are not affordable or reliable for the smallholder participants who represent the majority of the chain's actual nodes. It is this combination—fragmented records, an unverified ingestion step, and an accessibility gap at the most numerous and least resourced tier of the chain—that any credible transparency platform has to address together. The problem addressed in this paper is therefore to design a platform that preserves verifiable provenance across custody transfers, including batch splits and repackaging, while remaining simple and affordable enough for farm-gate and consumer use.

PROPOSED SYSTEM

A. Design Rationale

The gaps identified above motivate a platform built around three coupled capabilities that no single reviewed system combines: low-cost, farm-gate-accessible data capture that narrows the last-mile trust gap; a permissioned ledger that preserves verifiable provenance across batch splits and aggregation; and a single application layer that is simultaneously simple enough for a smallholder farmer to enter data into and for a consumer to query from, without either user needing to understand the underlying blockchain mechanics.

B. Expected Impact

Taken together, the low-cost sensor-assisted capture layer, the permissioned ledger with split-aware provenance tracking, and the dual-facing application layer are intended to compress the gap between when a custody event actually occurs and when it becomes verifiably, tamper-evidently recorded. Where a manual, paper-based traceability process might only be reconstructed after the fact during a recall investigation, taking days to trace a contaminated lot back to its source, a continuously updated ledger can in principle answer the same query in seconds, narrowing both the scope of a recall and the financial loss it causes. Over time, this shift from after-the-fact reconstruction to real-time, verifiable provenance is expected to reduce certification fraud, shorten payment-settlement times for

smallholder farmers, and give ordinary consumers a direct, evidence-based way to verify the claims made about the produce they buy, rather than relying on label text alone.

SYSTEM ARCHITECTURE

The proposed platform follows a three-tier architecture, shown in Figure 1. A data-capture layer deployed at each custody-transfer point—farm gate, aggregation center, cold-storage entry and exit, retail receipt—combines a low-cost RFID or QR tag applied to each batch or sub-batch with a lightweight mobile application that captures a timestamp, GPS location, and, where available, a sensor-verified attribute (weight, temperature, moisture) rather than a free-text field, so that ingestion defaults to an automatically captured signal and falls back to manual entry only when no sensor is available. A ledger layer implemented on Hyperledger Fabric records each custody-transfer event as a chaincode transaction within a consortium of farmers' cooperatives, aggregators, logistics providers, certifying bodies, and retailers, with each organization operating its own peer nodes and a practical Byzantine fault-tolerant ordering service coordinating transaction sequencing; bulky attachments such as certification documents or sensor-log files are stored off-chain on IPFS, with only their content hash and a retrieval pointer written on-chain, keeping ledger storage costs bounded while preserving cryptographic verifiability of the off-chain content.

Smart contracts (chaincode) encode the business rules identified as valuable in the reviewed literature: automatic payment release once a delivery-condition attestation is recorded, automatic certification-status propagation when a certifying body signs an approval transaction, and an explicit batch-split operation that creates linked child-lot identifiers preserving a verifiable reference back to the parent lot, directly addressing the batch-splitting gap noted in the Research Gap section.

On top of this ledger core, an application layer serves two distinct purposes through a shared backend but two distinct front ends. A farmer- and administrator-facing mobile application, designed for low-bandwidth and intermittent-connectivity conditions with offline queuing of capture events, guides a participant through a small number of structured prompts rather than a free-form form, and supports voice input in regional languages for participants with limited literacy. A consumer-facing mobile and web application allows a shopper to scan a retail package's QR code or RFID tag and retrieve a plain-language, chronologically ordered provenance summary—where a batch originated, how it was certified, how it was handled in transit—rendered from the same ledger without exposing raw transaction identifiers or requiring any blockchain-specific knowledge. This directly closes the gap between farmer-facing and consumer-facing interfaces built by different teams: both draw on the same ledger and the same backend query service.

Proposed Blockchain-Based Agricultural Supply-Chain Transparency Platform




Figure 1 Three-tier architecture of the proposed platform

METHODOLOGY / WORKFLOW

A. Operational Flow

A typical custody transfer begins when a participant—a farmer at harvest, an aggregator receiving a shipment, a retailer accepting delivery—scans or applies a tag and submits a capture event through the mobile application. The application forwards this event, together with any available sensor readings and GPS coordinates, to a backend gateway that constructs and submits the corresponding chaincode transaction to the Hyperledger Fabric network. Once endorsed by the required organizations' peers and committed by the ordering service, the transaction becomes part of the batch's permanent custody chain, and any associated bulk data is pinned to IPFS with its hash recorded on-chain. When a smart contract's specified condition is met—a quality attestation matching a pre-agreed threshold, for instance—the corresponding automated action (a payment release, a certification-status update) executes without further manual approval. At the consumer end, a QR-code scan triggers a read-only query against the ledger's current state and transaction history for that batch's identifier, which the backend translates into a plain-language summary before it reaches the consumer-facing application, closing the loop from a single farm-gate scan to a verifiable, consumer-readable provenance trail.

B. Challenges and Limitations

Although the proposed platform is designed to address the gaps identified in the reviewed literature, certain challenges and limitations remain. The present work is a literature review and system design; the architecture has not yet been validated against real, multi-season supply-chain data, and its fraud-detection and recall-response performance remain to be measured in partnership with an operating cooperative. The review itself is limited to eight sources, several of which are foundational surveys rather than production deployments.

At the architectural level, the platform requires internet connectivity at each ledger-writing node, and although farm-gate capture is designed to queue events offline, delays in synchronization can postpone the moment at which an event becomes verifiably recorded. Sensor-verified capture is available only where suitable low-cost sensing exists, so ingestion falls back to manual entry in other cases and the residual last-mile trust risk is reduced rather than eliminated. The correctness of smart-contract payment and certification logic depends on thresholds specified in advance, and a miscalibrated sensor can still cause an incorrect release or withholding of payment. The mapping between a physical tag and its on-chain record still depends on correct tag application at the packaging stage, and shipment-level in-transit monitoring adds hardware cost that must be balanced against its verification value.

CONCLUSION

The literature reviewed here shows steady, independent progress on the individual pieces of agricultural supply-chain transparency—immutable custody recording, smart-contract payment automation, cold-chain sensing, decentralized certification, and consumer-facing lookup—but no single existing system combines low-cost farm-gate verification, split-aware provenance tracking, and a genuinely accessible interface for both ends of the chain into one coherent platform. Each capability has so far been demonstrated in isolation, by a different team, under different assumptions about connectivity, hardware cost, and user technical skill, which makes it difficult for a cooperative or retailer to adopt more than one capability at a time without stitching together several unrelated systems.

The Blockchain-Based Agricultural Supply-Chain Transparency Platform proposed here is designed specifically to close that combination gap. A sensor-assisted, low-cost data-capture layer narrows the last-mile ingestion-trust gap, a Hyperledger Fabric ledger with explicit, split-aware batch-lineage support preserves verifiable provenance through the aggregation and repackaging steps that real agricultural handling routinely involves, and a dual-facing application layer—a low-literacy-friendly capture tool for farmers and a single-scan lookup tool for consumers—makes that verifiable provenance accessible to both ends of the chain from the same ledger. In doing so, the platform aims to shift agricultural supply chains away from opaque, trust-based coordination and toward transparent, evidence-based verification, reducing certification fraud, shortening payment cycles for smallholder

farmers, and giving consumers a direct, verifiable answer to where their food actually came from.

REFERENCES

- [1] Anonymous, "Blockchain-Enabled Farm-to-Fork Traceability Framework for Fresh Produce," unpublished survey.
- [2] Anonymous, "Smart Contract-Based Automated Payment and Quality Verification for Agricultural Trade," unpublished survey.
- [3] Anonymous, "A Survey on Consensus Mechanisms for Agricultural Supply Chain Blockchains," unpublished survey.
- [4] Anonymous, "IoT-Blockchain Integration for Cold-Chain Monitoring of Perishable Goods," unpublished survey.
- [5] Anonymous, "Decentralized Certification and Organic Produce Provenance Using Blockchain," unpublished survey.
- [6] Anonymous, "Hybrid Blockchain-Cloud Architecture for Multi-Stakeholder Agricultural Supply Chains," unpublished survey.
- [7] Anonymous, "QR-Code and RFID-Based Consumer-Facing Traceability Application," unpublished survey.
- [8] Anonymous, "Interoperability Challenges Across Blockchain Platforms in Cross-Border Agricultural Trade," unpublished survey.
- [9] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," , 2008.
- [10] V. Buterin, "Ethereum: A next-generation smart contract and decentralized application platform," , 2014.
- [11] E. Androulaki et al., "Hyperledger Fabric: A distributed operating system for permissioned blockchains," in *Proceedings of EuroSys*, 2018.
- [12] F. Tian, "An agri-food supply chain traceability system for China based on RFID & blockchain technology," in *Proceedings of the International Conference on Service Systems and Service Management (ICSSSM)*, 2016.
- [13] M. P. Caro, M. S. Ali, M. Vecchio, and R. Giaffreda, "Blockchain-based traceability in agri-food supply chain management: A practical implementation," in *Proceedings of the IoT Vertical and Topical Summit on Agriculture – Tuscany*, 2018.
- [14] E. D. Zamani et al., "Artificial intelligence and big data analytics for supply chain resilience: A systematic literature review," , 2022.
- [15] World Economic Forum, "Inclusive deployment of blockchain for supply chains," *White Paper Series*, 2019.
- [16] IBM Corporation, "IBM Food Trust: Blockchain for the world's food supply," *Technical Overview*, 2020.