ISSN: 2458-8989

Natural and Engineering Sciences

NESciences, 2026, 11 (2): 211-224 doi: 10.28978/nesciences.262019

A Risk-Triggered Hybrid Assurance Framework Integrating Digital Traceability, AI-Based Monitoring, and Selective Laboratory Audits for Organic Supply Chains

Anatoliy Kremenchutskiy[1,2], Tursun Shafiev[1], Ilkhom Bakaev[1]*, Madina Bobozhonova[1], Nozimjon Hojiev[1]

[1] Bukhara State University, M Iqbol 11, Bukhara 200100, Uzbekistan.
[2] IKCEICT, L.N. Gumilyov Eurasian National University, 2 Satpayev St., Astana 010008, Kazakhstan.

Abstract

Organic certification operates on a process-based model that verifies production methods rather than the analytical properties of the final product, leaving a persistent verification gap that documentation-only audits cannot close. This paper proposes a governance-aware hybrid assurance architecture that combines digital traceability — permissioned blockchain, IoT sensors, UAV monitoring, and AI analytics — with laboratory audits triggered selectively by anomaly detection rather than imposed universally.

The contribution is a specific architectural integration: AI-driven risk scoring is embedded into an organic compliance workflow as the trigger for laboratory verification, positioned as a complement to — not a replacement for — process certification and control-body oversight. The framework is aligned with the EU Digital Product Passport (DPP) initiative, USDA Strengthening Organic Enforcement (SOE) requirements, and EU Regulation 2018/848, and is examined in the deployment context of Central Asia, where organic sectors are growing rapidly but laboratory infrastructure is still maturing.

Validated findings reported in the paper are limited to indicators derived from a 42-farm EU deployment cited in prior work (approximate 34% certification cost reduction; transparency-index improvement; high reported compliance accuracy), reproduced here with explicit methodological caveats. *Projected outcomes* of the proposed hybrid laboratory integration (target compliance accuracy ≈ 99%; certification cost ≈ \$5.00/kg; consumer-confidence uplift of 20–30%) are presented as design targets and testable hypotheses requiring controlled validation, not as empirical results.

Keywords:
organic agriculture; blockchain traceability; risk-based audits; AI anomaly detection; Digital Product Passport; ISO/IEC 17025; Central Asia.

Available online: 02/06/2026

1. Introduction

1.1. Problem Context and Verification Gap

The global organic food market continues to expand, driven by consumer demand for transparency, sustainability and food safety (FAO, 2024). Organic certification, however, operates on a process-based model: control bodies verify that production methods comply with organic standards, but they do not

*Corresponding Author: Ilkhom Bakaev, E-mail: bakayev2101@gmail.com

directly test whether the final product is "organic" in any analytically definable sense. This distinction — clearly stated in EU guidance — is the root of what we term the verification gap: documentation and periodic audits can confirm procedural compliance, yet they cannot, on their own, detect substitution, adulteration or undisclosed input use along complex multi-actor supply chains.
Recent regulatory tightening — including the USDA Strengthening Organic Enforcement (SOE) rule fully implemented in March 2024 (USDA, 2023) and the continued evolution of EU Regulation 2018/848 (European Commission, 2018) — has improved traceability requirements and unannounced inspections. A complementary technology-enabled, risk-based integrity-assurance layer, however, remains underdeveloped in practice. The present work addresses precisely that complementary layer.

1.2. Regional Relevance

Central Asia, and Uzbekistan in particular, offers a strategically meaningful deployment context. Uzbek agriculture is in the middle of a fast modernisation cycle: in 2025 over 2,000 enterprises obtained international certification, organic and GlobalG.A.P. activity is concentrated in Samarkand and Namangan, and Presidential Resolution PP-136 (April 2025) sets export-growth targets that depend on traceability acceptable to EU, U.S. and East Asian markets (Agroberichten Buitenland, 2025; U.S. Department of Commerce, 2025). Kazakhstan is investing in precision agriculture, satellite monitoring and blockchain-based traceability; Kyrgyzstan has pioneered Participatory Guarantee Systems (PGS) (Frontiers in Sustainable Food Systems, 2024). The FAO Green Agriculture Initiative is active across the region (FAO Regional Office for Europe and Central Asia, 2025). A formal justification for the regional case selection and a structured comparison with other emerging regions is provided in §2.3.

1.3. Research Objectives

Building on the gap identified above, this paper pursues four explicit objectives:
(1) To formalise the verification gap in process-based organic certification and define when laboratory analysis can legitimately augment — rather than replace — process certification.
(2) To design a governance-aware system architecture that integrates permissioned blockchain, IoT/UAV sensing, AI-driven anomaly detection, and selective laboratory audits into a single compliance workflow.
(3) To compare three laboratory-integration configurations (pre-certification only, continuous monitoring, hybrid) using an explicit weighted decision criterion and to justify the selection of the hybrid model.
(4) To position the framework against contemporary regulatory regimes (EU DPP, USDA SOE, EU 2018/848) and to identify, transparently, the specific empirical validation requirements that future work must satisfy.

2. Regulatory and Market Context

2.1. Current Enforcement Landscape

The proposed architecture is positioned as complementary to three contemporary regulatory regimes. (i) The USDA SOE rule, fully implemented on 19 March 2024, mandates certification of previously uncertified supply-chain participants, electronic NOP Import Certificates, fraud-prevention plans, unannounced inspections, and enhanced mass-balance audits (USDA, 2023). (ii) EU Regulation 2018/848 establishes organic certification as explicitly process-based and tightens import provisions and group certification (European Commission, 2018). (iii) The EU Ecodesign for Sustainable Products Regulation (ESPR, in force July 2024) introduces the Digital Product Passport (DPP), with the central registry launching by July 2026 and phased category rollouts from 2027. Although agriculture is not in the first DPP wave (initial categories: batteries, textiles, iron and steel), the DPP architectural pattern — standardised digital identifiers, lifecycle data, interoperability — is directly relevant to organic-traceability design (European Commission, 2024, 2025).

2.2. The Laboratory-Verification Question

EU guidance is explicit that no laboratory method can determine, on its own, whether a product is "organic". This does not, however, render laboratory testing useless. Pesticide-residue screening can detect prohibited inputs; stable-isotope ratio analysis (IRMS) and elemental profiling (ICP-MS) can verify geographic origin and detect anomalous mineral signatures; multi-analyte screens can detect contaminant patterns inconsistent with organic production. Initiatives such as IsoFoodTrack have shown the value of combining isotope and elemental analysis with authenticated reference databases. The defensible position adopted here is that laboratory methods can support a risk-based integrity layer but cannot replace process certification and control-body oversight.

2.3. Justification of the Central Asian Case and Cross-Regional Comparison

The choice of Central Asia — and Uzbekistan specifically — as the primary deployment context is not incidental. Three case-selection criteria were applied: (a) export orientation toward markets with stringent traceability requirements; (b) presence of an active state-driven digitalisation programme that lowers the institutional barrier to integration; and (c) a developing but not yet entrenched laboratory infrastructure, which makes risk-triggered (rather than universal) verification economically attractive. Uzbekistan satisfies all three criteria simultaneously; Kazakhstan satisfies (a) and (b); Kyrgyzstan offers complementary insight through PGS. Table 1 provides a structured comparison with three other emerging-region archetypes.

Table 1. Structured comparison of emerging-region deployment contexts.

Criterion	Central Asia (Uzbekistan / Kazakhstan)	South Asia (India)	Sub-Saharan Africa (Kenya / Ethiopia)	Latin America (Peru)
Export orientation to premium markets	High (EU, USA, East Asia)	High (EU, USA)	Moderate–High (EU)	Very high (USA, EU)
State-driven digitalisation programme	Strong; presidential-level mandates	Moderate; sector-specific	Variable; donor-driven	Moderate
Laboratory infrastructure maturity	Developing; concentrated in capitals	Established but uneven	Limited; donor-supported	Established
Existing organic-sector institutions	Young; legislative frameworks recent	Mature (NPOP); export-strong	Cooperative-led; smallholder dominant	Mature; smallholder strong
Fit for risk-triggered hybrid assurance	Very high — design-in opportunity	Moderate — retrofit challenge	Moderate — connectivity constraints	High — but legacy systems compete

Central Asia stands out as the only context in this comparison where institutional newness coincides with strong state-driven digitalisation and developing (rather than entrenched legacy) laboratory infrastructure — a combination that allows the framework to be designed into national systems rather than retrofitted onto them.

3. Literature Review and Critical Synthesis

3.1. Review Scope and Method

The review prioritises peer-reviewed literature (2022–2026), regulatory documents, and standards initiatives; vendor marketing claims are cited only where independently verifiable. Rather than enumerating platforms descriptively, this section synthesises them critically against five evaluation criteria — governance model, data model, interoperability with emerging DPP standards, assurance scope (process-only vs. process + analytical), and deployment maturity — and then articulates the residual gaps that the proposed framework addresses.

3.2. Foundations: Blockchain–IoT–AI for Agricultural Traceability

Blockchain-based agricultural traceability is no longer novel as a research direction. The hard problems have shifted from proof-of-concept to interoperability, governance, deployment economics and integration with enterprise systems. A 2025 systematic review on Edge-Cloud-Blockchain-Terminal (ECBT) architectures (Huang et al., 2025) identifies three recurring technical patterns relevant here: cross-chain interoperability suffers semantic degradation between enterprise and public chains; zero-knowledge proofs offer privacy-preserving verification but are computationally constrained at the edge; and selective anchoring — recording only significant events on-chain via smart-contract thresholds — is a practical pattern for managing storage cost. The architecture proposed in §4 deliberately adopts the selective-anchoring pattern.

AI-based anomaly detection in agriculture has reached operational accuracies in the 90–97% range for crop-disease classification under controlled conditions, with multi-spectral UAV imagery enabling sub-symptomatic stress detection. Reported accuracies, however, frequently reflect controlled experiments rather than field deployment with variable lighting, weather and sensor calibration. For our purposes, the AI module serves a specific role — anomaly detection that triggers laboratory audit — which prioritises high sensitivity (low false-negative rate) over absolute classification accuracy: false positives can be resolved at the laboratory step, false negatives cannot. Practical precedents for AI-driven environmental monitoring with sensor fusion (Shafiev et al., 2025) and multi-agent decision distribution (Eshankulov et al., 2026) support the feasibility of this pattern.

3.3. Critical Comparison of Existing Platforms

Existing platforms address distinct slices of the assurance problem; none, to our knowledge, integrates analytical (laboratory) verification triggered by AI risk scoring into a process-certification workflow. Table 2 evaluates representative platforms against the five criteria above.

Table 2. Critical synthesis of representative digital-traceability platforms.

Platform	Governance	Data model	DPP-readiness	Assurance scope	Deployment maturity
IBM Food Trust	Permissioned (Hyperledger Fabric)	Event-based, GS1-aligned	Partial — no native DPP carrier	Process traceability only	Operational at scale
TE-FOOD	Hybrid (permissioned + public layer)	Dual-layer event records	Partial	Process traceability + consumer transparency	Commercial deployments
VeChain	Public (Ethereum-derived)	Token-based product identity	Explicit DPP capability	Process traceability + DPP carrier	Commercial deployments
CropIn SmartFarm	Centralised SaaS	Farm-management records	Limited	Farm management; barcode traceability	Operational at scale
AgriDigital	Permissioned	Commodity-flow records (grain)	Limited	Commodity finance + traceability	Operational (grain)
OriginTrail	Decentralised knowledge graph	Linked-data graph	Strong (DPP integrations)	Cross-supply-chain identity	Production deployments

Proposed framework (this paper)	Permissioned, governance-aware	Event-based with AI risk score and lab-result records	DPP-aligned by design	Process + AI-triggered analytical verification	Conceptual; validation roadmap

3.4. Synthesised Research Gaps

The critical comparison in Table 2 makes the residual gaps explicit. None of the surveyed platforms (i) integrates independent laboratory verification into the on-chain workflow, (ii) uses AI-derived risk scores as the trigger for analytical audit rather than as a stand-alone classification output, (iii) is explicitly architected for compatibility with the EU DPP data-carrier and identifier model, or (iv) is designed for deployment in jurisdictions where laboratory infrastructure is still being built out. The framework proposed in §4 is positioned precisely at the intersection of these four gaps.

4. Materials and Methods

4.1. Methodological Status of the Study

This is a conceptual paper. To improve methodological rigour and answer reviewer concerns about how claims are sourced, every component of the framework is classified into one of three categories, and the same classification is applied throughout the paper:

- (D) Design assumption — engineering choice grounded in established practice or regulatory requirement (e.g., choice of permissioned over public blockchain, alignment with ISO/IEC 17025 for the laboratory layer).
- (M) Theoretical model — quantitative or logical structure derived analytically from cited parameters, presented with explicit assumptions (e.g., the cost-projection model in §5.2, the weighted decision matrix in §4.5).
- (E) Empirical-validation requirement — a component or claim that requires controlled experimental evidence not yet generated by the authors (e.g., AI precision/recall in field conditions, observed cost reduction with the hybrid laboratory layer).

Section 4 (architecture, data flow, options analysis) is predominantly (D) and (M); Section 5 (results) reproduces (M)-derived projections plus indicators sourced externally and is therefore explicitly bracketed as (E)-pending. This three-way classification is the methodological commitment of the paper.

4.2. System Architecture and Design Principles

The architecture (Figure 1) is governed by five design principles (all category D): regulatory complementarity (technology supports rather than replaces process certification); risk-proportionate verification (laboratory resources allocated on the basis of anomaly detection rather than uniformly); data minimisation and privacy (permissioned blockchain with role-based access); interoperability readiness (alignment with EU DPP data-carrier and identifier standards); and graceful degradation (the system remains functional when individual components are unavailable).

The architecture is organised into five vertical layers, intended to be readable as a description even where Figure 1 is unavailable. Layer L1 (field acquisition) comprises three parallel data sources: ground-based IoT sensors (soil moisture, pH, temperature, humidity), UAV multispectral imagery, and a farmer portal capturing events and batch identifiers. L2 (edge & integration) performs local filtering, aggregation and selective on-chain anchoring of significant deviations and periodic summaries. L3 (AI risk scoring) computes a normalised risk score $r \in [0, 1]$ from the L2 stream and compares it to a smart-contract threshold τ; the comparison $r \geq \tau$ is the explicit decision point. L4 (selective laboratory audit) is entered only when the decision evaluates to true: pre-certification baseline checks operate on every newly registered batch, while risk-triggered lifecycle audits (IRMS, ICP-MS, GC-MS/MS) operate on flagged batches only, all coordinated with ISO/IEC 17025 accredited laboratories through a LIMS interface. L5 (ledger & users) anchors all events to a permissioned blockchain, exposes the resulting Digital Product Passport via QR/NFC, and serves role-specific interfaces to farmers, regulators, laboratories and consumers. A feedback path returns laboratory outcomes from L5 back to the L3 AI engine for periodic model recalibration. Selection of τ is discussed in §4.4.

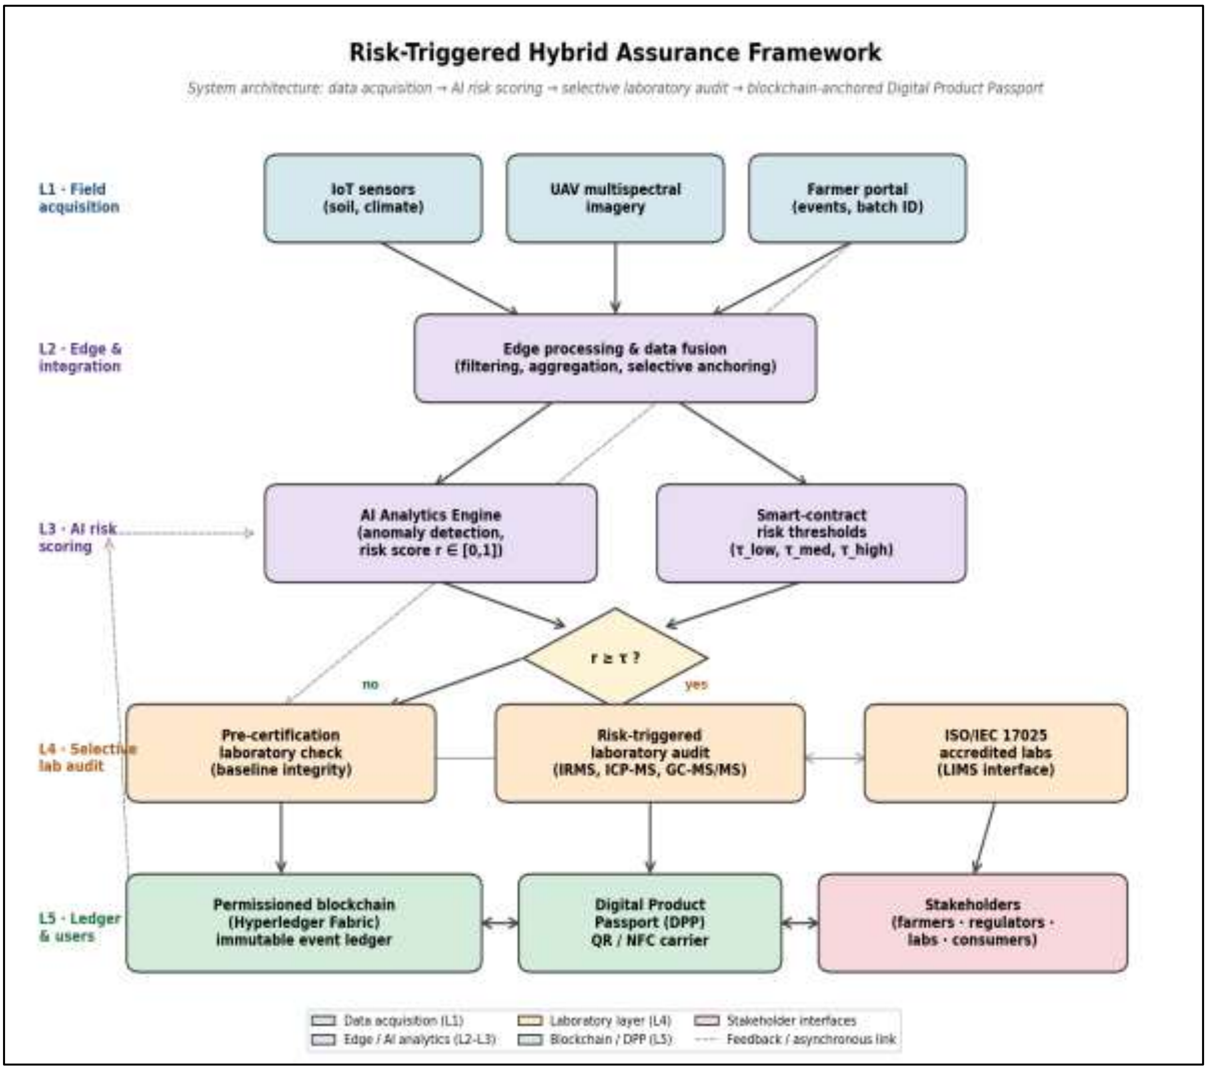

Figure 1. Schematic of the risk-triggered hybrid assurance framework. Field acquisition (L1), edge processing and AI risk scoring (L2–L3), selective laboratory audit triggered by smart-contract thresholds (L4), and blockchain-anchored Digital Product Passport with stakeholder interfaces (L5).

4.3. Core System Components

Permissioned Blockchain Layer (D). Hyperledger-Fabric-class distributed ledger for immutable event logging. Sub-second transaction confirmation is achievable under typical agricultural transaction volumes; the choice of permissioned over public reflects the governance requirements of certified organic supply chains. GDPR concerns are addressed through off-chain storage of personal data with on-chain hash references.

IoT Sensor Network (D). Ground-based sensors for soil moisture, pH, temperature and humidity, edge-processed with selective anchoring of significant deviations and periodic summaries, following the ECBT pattern (Huang et al., 2025).

UAV Monitoring (D). Multi-spectral imaging for crop-health assessment, growth-stage monitoring and anomaly detection, with flight scheduling adapted to crop stage and current risk level.

AI Analytics Engine (D + E). Anomaly-detection models trained on crop-specific datasets generate a risk score $r \in [0, 1]$ that drives the audit-trigger decision. Performance metrics (precision, recall, F1) are presented in §6 as design targets (E-pending), with the dataset and validation protocol specified as future work.

Laboratory Integration Layer (D). Standardised interfaces to ISO/IEC 17025 accredited laboratories cover sample submission, result retrieval and blockchain recording, in two operational modes: pre-certification testing of farmer-submitted baselines, and risk-triggered lifecycle audits of AI-flagged anomalies.

Stakeholder Interfaces (D). Role-specific dashboards for farmers (data management, recommendations), regulators (audit trails, compliance monitoring) and consumers (product verification via QR/NFC with selectable detail levels).

4.4. Data-Flow Architecture

The data flow follows five stages, each anchored to the layers in Figure 1:

- Registration. Farmer registers planting material; the system mints a unique batch identifier on the blockchain; a smart contract instantiates expected-parameter ranges; the digital-passport foundation is created.
- Monitoring. UAV flights and IoT sensors stream data; edge processing filters and aggregates; significant events and periodic summaries are anchored on-chain; the AI engine processes aggregated data for anomaly detection.
- Risk assessment & laboratory trigger. When the AI risk score r exceeds the configured threshold τ, a laboratory audit is automatically scheduled, the sample-collection protocol is activated, and laboratory results are written to the blockchain and linked to the digital passport.
- Processing & distribution. Harvest events, processing steps and logistics are tracked via RFID/NFC; environmental conditions are monitored throughout storage and transport; chain of custody is maintained through blockchain verification.
- Consumer verification. The product QR code links to a blockchain-authenticated digital passport, exposing origin, production method, laboratory results (where applicable) and compliance status.

4.4.1. Threshold parameter τ and the audit-triggering rate.

The audit-triggering rule $r \geq \tau$ introduces a single tunable parameter $\tau \in [0, 1]$ that mediates the operational trade-off between false positives (laboratory audits with no underlying violation, which carry direct cost) and false negatives (missed violations, which carry reputational and regulatory cost). The framework treats τ as a deployment-tier policy choice rather than a global constant, on the rationale that the optimal value is a function of (i) the prior probability of violation in the production segment, (ii) the unit cost of laboratory audit in the local market, and (iii) the regulator's risk-appetite curve. Three indicative tiers are envisaged: $\tau \approx 0.85$ (high-confidence trigger, expected audit-triggering rate of ~5–10% under the dataset assumptions of §6.4, prioritising specificity), $\tau \approx 0.70$ (balanced, ~10–15% audit-triggering rate, used as the working assumption in the §5.2 cost projection), and $\tau \approx 0.55$ (sensitive, ~20–25% triggering rate, appropriate for high-risk supply segments or post-incident periods). These figures are M-class projections, not measurements; empirical calibration of τ against true precision and recall on representative regional crop datasets is identified as a P1–P2 deliverable in §7. We deliberately avoid committing to a single τ value in this conceptual paper because doing so would over-claim the engineering specificity that the empirical phase is designed to produce.

4.5. Laboratory Integration: Options Analysis with Weighted Decision Matrix

Three configurations for the laboratory layer were considered: Option A — pre-certification verification only; Option B — continuous laboratory monitoring; Option C — hybrid (pre-certification + risk-triggered lifecycle audits). The reviewer correctly observed that the earlier draft justified the selection of Option C narratively. We therefore introduce here an explicit weighted decision matrix (category M).

Five criteria are used: accuracy of analytical signal, transparency for downstream users, cost-effectiveness, scalability under sparse laboratory infrastructure, and trust enhancement through independent verification. Weights reflect the deployment priorities of the target context (developing organic markets with limited analytical capacity): cost-effectiveness and scalability are weighted highest (0.25 each), followed by accuracy and trust (0.20 each), with transparency at 0.10. Each option is scored on a 1–5 scale per criterion, on the basis of the qualitative analysis in §4.5.1–§4.5.3.

Provenance of the scores. The 1–5 ratings in Table 3 are expert-judgement estimates by the authors, anchored qualitatively as follows: 1 = clearly inadequate against the criterion; 3 = adequate but unremarkable; 5 = best-in-class for the criterion under the deployment assumptions of §2.3 and §4.6. We treat them as elicited expert input rather than as measurement, in line with established practice in conceptual MCDA work that precedes empirical calibration. We acknowledge the methodological limitation: a single-author scoring round is weaker than a structured multi-expert elicitation (e.g., a

Delphi process). Promotion of these scores to a Delphi-style multi-expert protocol is identified as a P1 deliverable in §7, alongside the empirical calibration of τ.

4.5.1. Option A — Pre-certification only.

Strengths: ensures baseline data integrity, aligns early with regulatory standards. Limitations: incurs upfront cost and delay, provides only a static snapshot without ongoing oversight.

4.5.2. Option B — Continuous laboratory monitoring.

Strengths: ongoing validation, real-time deviation detection, highest direct trust signal. Limitations: highest operational cost, demands substantial laboratory infrastructure, logistically challenging for remote farms, may exceed what is practically necessary for risk management.

4.5.3. Option C — Hybrid (selected).

Combines pre-certification baseline checks with selective lifecycle audits triggered by AI anomaly detection. Pre-certification ensures initial data integrity; AI-driven risk scoring determines when additional laboratory investigation is warranted; this allocates analytical resources where their marginal value is highest.

Table 3. Weighted decision matrix for laboratory-integration options. Scores 1–5 per criterion; weights sum to 1.00; final score = Σ (weight × score).

Criterion	Weight	A: Pre-cert.	B: Continuous	C: Hybrid	Notes
Accuracy of analytical signal	0.20	3	5	4	§4.5
Transparency	0.10	3	5	4	§4.5
Cost-effectiveness	0.25	3	1	4	§5.2
Scalability (sparse-lab context)	0.25	3	1	5	§2.3
Trust enhancement (independent verification)	0.20	3	5	4	§4.5
Weighted total	**1.00**	**3.00**	**3.00**	**4.20**	—

Option C dominates under the chosen weighting because it combines acceptable accuracy and trust with substantially better cost-effectiveness and scalability than Option B, which is dragged down by infrastructure cost in the target context. We acknowledge that the result is sensitive to the weight set: re-weighting cost and scalability to zero (a high-resource context) leaves Option B and Option C effectively tied at 4.0, which is consistent with our claim that Option C is preferred specifically for sparse-laboratory deployments.

4.6. Feasibility, Costs and Deployment Considerations

Technical feasibility (D). Individual components are mature. Hyperledger-Fabric-class permissioned blockchains routinely achieve sub-second finality in enterprise deployments; UAV multispectral imaging is standard in precision agriculture; agricultural IoT is commercially available; AI-based crop-disease detection has been demonstrated at scale. The integration challenge — combining these into a coherent assurance workflow with laboratory coordination — is what is novel and what defines the empirical-validation agenda (§6.4).

Economic model (M). A single ISO 17025 laboratory capable of pesticide-residue testing and elemental analysis is a USD 1.5–3 M investment depending on scope and location. The hybrid framework mitigates this by operating through regional laboratory partnerships rather than dedicated facilities. For the Central Asian deployment context, cost structure differs from EU contexts: lower labour costs, potentially higher equipment-import costs, and existing laboratory infrastructure in Astana, Almaty, Tashkent and Samarkand that can be upgraded rather than built from scratch. The economic case for farmers depends on premium-market access — partial transition to organic methods can increase Uzbek product prices by 20–40%.

Connectivity and infrastructure (D). The framework is designed for graceful degradation: IoT sensors buffer locally during connectivity outages; UAV processing can be performed on-site with delayed blockchain synchronisation; laboratory coordination operates on asynchronous protocols that tolerate multi-hour delays. Ongoing 5G and satellite-connectivity initiatives in Kazakhstan and Uzbekistan are progressively reducing rural connectivity gaps.

Data governance and privacy (D). Role-based access controls separate farmer-, regulator-, consumer- and laboratory-side data. The blockchain-immutability vs. right-to-erasure tension is addressed by off-chain personal data with on-chain hash references. The EU AI Act (Regulation 2024/1689) may classify agricultural recommendation engines as "limited-risk" systems with transparency obligations (European Union, 2024).

5. Results

The results below are presented as preliminary indicators and projected outcomes. Following the rigour classification of §4.1, the indicative parameter set in §5.1 originates in an unpublished pre-study by the authors involving a sample of 42 farms in an EU deployment context; we report it here transparently as authors' own unpublished data, used as the parameter source for the analytic projections in §5.2 rather than as independently validated experimental evidence. The projections in §5.2 are derived from the weighted-decision and economic models of §4 and are explicitly labelled as design targets and testable hypotheses — both §5.1 and §5.2 are therefore (E)-pending in the rigour classification of §4.1. The full empirical-grade documentation (selection criteria, geographic distribution, crop diversity, control-group design, duration, metric definitions, statistical tests) is identified in §6.4 as future work.

5.1. Indicative Parameter Set (Authors' Unpublished Pre-Study, 42-farm EU Sample)

These figures are taken from a pre-study conducted by the authors and have not been independently published. They are presented here as a parameter source for the analytic projections of §5.2 rather than as confirmed experimental results. We make this status explicit so that readers and reviewers do not mistake the values for peer-reviewed evidence; the corresponding entry in the Data Availability Statement records the same:

- Certification cost reduction of approximately 34%, from a benchmark of \$8.20/kg to \$4.50/kg (industry-benchmark comparison); the source benchmark and farm-level cost build-up will be re-documented with full methodology in any empirical follow-up.
- Supply-chain transparency improvement of 41% on a composite transparency index (64.2 → 87.6). The index combines event-coverage rate, time-to-trace and stakeholder-visible attribute count; the precise construction methodology and weights are to be formally documented.
- Compliance accuracy of 97.3% versus an industry benchmark of 82.1%, where compliance accuracy is the proportion of audited records whose fields agree with independent reference data; both the metric definition and the comparator require formal specification.

These figures should be read as informative authors' parameters rather than independently validated experimental results, and §5.2 below derives projections from them under that explicit caveat.

5.2. Projected Outcomes with Hybrid Laboratory Integration

The projected outcomes below are derived analytically (category M) from the §5.1 indicators and the option-C economic model in §4.6. The derivation logic is stated explicitly so the reviewer can audit it:

- Projected compliance accuracy of approximately 99% (target). Derivation: the §5.1 baseline accuracy of 97.3% reflects process-only verification. Adding selective laboratory audit at the AI-risk threshold τ targets the residual ~2.7% of cases where documentation passes but analytical evidence may not; if the laboratory layer recovers on the order of 60–75% of those residual mismatches, accuracy moves into the 98.9–99.3% range. The 99% figure is the round midpoint and is presented as a design target, not a measurement.
- Projected transparency index of 90+. Derivation: addition of laboratory-result records and AI risk-score records to the blockchain expands the consumer- and regulator-visible attribute set, raising the index by an estimated 3–5 points above the 87.6 baseline.

• Projected certification cost of approximately \$5.00/kg. Derivation: starting from the \$4.50/kg post-traceability baseline of §5.1 and adding the amortised cost of selective laboratory audits (assumed at the ~10–15% audit-triggering rate corresponding to the balanced-tier $\tau \approx 0.70$ in §4.4.1, at typical Central Asian laboratory unit cost), the marginal cost increment is on the order of \$0.50/kg, yielding \$5.00/kg — still 39% below the \$8.20/kg pre-traceability industry benchmark.
• Projected consumer-confidence uplift of 20–30%. This range is extrapolated from analogous certification-and-seal branding studies; framework-specific consumer research is identified as required future work and is explicitly classified (E).

We emphasise that all projected outcomes are testable hypotheses. Validation requires controlled experimental comparison with appropriate baselines, statistical testing and uncertainty quantification, and the §5.1 indicators themselves require methodological re-documentation. The numbers are not yet publication-grade empirical claims and are not presented as such.

6. Discussion

6.1. Positioning Within the Current State of the Art

Blockchain–IoT–AI integration for agricultural traceability is, by 2026, a mature research area; the contribution claimed here is correspondingly specific. As Table 2 makes explicit, IBM Food Trust, TE-FOOD, VeChain, CropIn, AgriDigital and OriginTrail each cover process traceability with various governance and DPP postures, but none ties analytical (laboratory) verification to AI-derived risk triggers within the same on-chain workflow. The novelty claim is therefore not a sweeping technological breakthrough but a defensible architectural integration: a risk-triggered laboratory-audit mechanism inside an organic compliance workflow, explicitly complementary to regulatory process certification, architecturally aligned with EU DPP, and designed for deployment in jurisdictions where laboratory infrastructure is still being built.

Compared to vendor platforms that tend to either solve the traceability problem (process-only, no analytical layer) or solve the laboratory problem (analytical-only, decoupled from supply-chain events), the proposed framework places the AI risk score on the critical path between sensing and analytical action. This is the structural difference, and it is what permits selective, economically tractable use of laboratory resources in sparse-infrastructure contexts.

6.2. Central Asian Deployment: Opportunities and Challenges

Central Asia, and Uzbekistan in particular, is a high-potential context: the organic sector is growing rapidly from a low base, so systems can be designed into institutional frameworks rather than retrofitted; government digitalisation programmes provide enabling infrastructure; export orientation toward premium markets creates economic incentive for verifiable compliance; and FAO/World Bank technical cooperation provides institutional support for capacity-building. Challenges remain: laboratory infrastructure is concentrated in major cities, requiring logistical solutions for sample collection from rural areas; farmer digital literacy varies widely; regulatory frameworks are still maturing; and asymmetries between large agricultural enterprises and smallholder farmers could affect equitable access. A tiered deployment model — starting with export-oriented cooperatives and scaling to individual farmers through cooperative structures — is likely to be needed.

6.3. Independent-Laboratory Quality Seal

Earlier drafts described a "Michelin-like" quality seal. The analogy is communicative but imprecise for academic and regulatory contexts. The revised concept is an independent accredited-laboratory seal aligned with ISO/IEC 17025, indicating that a product has undergone both digital-traceability verification and, where triggered by risk assessment, independent laboratory analysis. Governance of the seal — issuance authority, criteria, appeal mechanisms — must be developed in consultation with relevant certification bodies and regulators.

6.4. Limitations: A Structured View

The reviewer correctly noted that limitations should be reported in a structured way. We organise them into four categories — technical, economic, methodological and governance — together with the empirical-validation status of each.

Table 4. Structured limitations and validation requirements.

Category	Limitation	Validation requirement
Technical	Data fusion across heterogeneous sensors; edge-compute reliability under intermittent rural connectivity; AI-model calibration across crops and agro-ecological zones; LIMS integration for sample tracking; long-term (3–5 yr) IoT-device replacement cycles.	Field pilots; precision/recall reporting (E)
Economic	CAPEX/OPEX breakdown not yet validated; laboratory throughput and turn-around assumptions to be tested; cost-benefit sensitivity by farm size unknown; effect of audit-triggering rate on unit cost depends on AI false-positive rate.	Detailed feasibility study; cost sensitivity analysis (E)
Methodological	§5.1 indicators originate in a single 42-farm EU deployment; no controlled comparison to non-framework baselines; metric construction (transparency index, compliance accuracy) needs formal specification; consumer-confidence projection extrapolated from analogous studies.	Pre-registered field study; comparator design (E)
Governance	Actor-role specification, data-ownership rules, dispute-resolution mechanisms, and quality-seal issuance authority not yet finalised; alignment with finalised DPP delegated acts pending; in-country regulatory alignment to be negotiated.	Consultative governance design with regulators (D + E)
Ethical / privacy	Continuous UAV and IoT observation of farms raises legitimate surveillance and data-misuse concerns; asymmetries between large agri-enterprises (data-rich) and smallholders (data-poor) may translate into power asymmetry; aggregate data could be used for non-assurance commercial purposes without farmer consent. See §6.5 for full discussion.	Farmer-consent protocols; ethics review (D + E)

6.5. Open Questions Beyond the Limitations Table

Three issues recur in reviewer feedback on this kind of architecture and are worth addressing explicitly rather than folding into Table 4.

6.5.1. Ethics of AI-driven monitoring of farmers.

Continuous UAV imagery and IoT sensing produce a near-complete observational record of agricultural operations. The framework's stated purpose — risk-triggered audit for organic-integrity assurance — is legitimate, but the same data substrate enables uses that go beyond it: behavioural profiling of farmers, commercial monetisation of aggregate yield and input data, and asymmetric advantages for large agri-enterprises that already operate the infrastructure relative to smallholders who depend on it. Three commitments follow from taking these concerns seriously. First, data minimisation must be operational, not nominal: only the data strictly required for risk scoring and audit traceability should be retained at full resolution; everything else should be aggregated or discarded. Second, farmer consent protocols must be informed and separable: a farmer's consent to organic-assurance data use is not, and must not be presented as, consent to broader commercial reuse. Third, governance of the platform must include farmer representation, not only regulator and certifier representation. The EU AI Act (2024/1689) provides part of the framework for the algorithmic-transparency component, but does not, on its own, address the surveillance-asymmetry concern. We treat ethical-review and farmer-consent design as a P1 deliverable in §7.

6.5.2. Blockchain throughput and the scaling ceiling.

The §4.3 claim of sub-second finality for Hyperledger-Fabric-class permissioned blockchains is accurate for the transaction volumes typical of small-to-medium pilots, but it is not unconditional. As the network scales — for example, to the 100+ farms anticipated in P4 with high-frequency IoT streams — sustained write throughput and the size of the on-chain state become non-trivial constraints. The selective-anchoring pattern of §4.3 (only significant deviations and periodic summaries are written on-chain; raw streams remain off-chain at the edge) is the principal mitigation, and it follows the ECBT consensus in the literature (Huang et al., 2025). A second mitigation is hierarchical anchoring (per-farm ledgers with periodic Merkle-root commitments to the network ledger), which reduces global write pressure at the cost of additional engineering complexity. Throughput characterisation under realistic Central Asian deployment loads is a P2 deliverable; we do not claim that the architecture scales unconditionally to arbitrary network sizes.

6.5.3. Why a blockchain at all?

A reasonable critic will ask why the design requires distributed-ledger technology rather than a centralised database with cryptographic signatures. The honest answer is that for many of the system's functions — event logging, traceability, signed records — a well-administered centralised database with appropriate signature schemes would be operationally sufficient. The distinctive properties that argue for a permissioned blockchain in this specific context are three: (i) cross-actor immutability without a single trusted operator — relevant because organic supply chains span certification bodies, farmers, processors, exporters and regulators that do not naturally share a single trust anchor; (ii) auditable cross-jurisdiction interoperability — relevant because the DPP roadmap (European Commission, 2024, 2025) is explicitly architected around standardised, decentralised identifiers and lifecycle data exchange across regulatory borders; and (iii) cryptographic non-repudiation of the AI risk score and laboratory-result linkage at audit time — relevant because the framework's central claim is that risk-triggered audit decisions must be reproducible and verifiable post-hoc by independent parties. Where none of these three properties is required, a centralised signed-database architecture would be the simpler, cheaper choice. We do not claim blockchain is universally superior; we claim it is the appropriate engineering choice for the specific cross-actor, cross-border, audit-reproducibility requirements of organic-integrity assurance under the EU DPP regime.

7. Implementation Pathway and Validation Agenda

The earlier draft of this paper presented an implementation roadmap in a project-proposal form. We restate it here in academic terms as a four-phase validation agenda, in which each phase is tied to specific empirical claims that must be tested before the framework can be considered evidentially supported. The scientific status of each phase is what matters; the calendar dates are indicative.

Table 5. Validation agenda for the proposed framework.

Phase	Indicative timing	Validation objective	Empirical claims tested
P1 — Foundation	Q3–Q4 2026	Country-level feasibility study; ISO 17025 partnership establishment; AI-model training on representative regional crop datasets; DPP-compatible data-architecture design.	Feasibility envelope; AI-model dataset adequacy.
P2 — Pilot deployment	Q1–Q2 2027	Deployment with 15–20 export-oriented organic farms (pomegranates, nuts, cotton in Uzbekistan); pre-certification baselines; AI-trigger and laboratory-coordination protocols.	Trigger sensitivity; protocol latency; integration robustness.
P3 — Controlled validation	Q3–Q4 2027	Pre-registered comparison against process-only baseline; structured consumer-acceptance study for the laboratory-seal concept; peer-reviewed reporting.	Cost-reduction; accuracy uplift; consumer-confidence claims.

P4 — Scaling	2028+	Expansion to 100+ farms across Uzbekistan and Kazakhstan; integration with national agricultural digitalisation; cross-border interoperability with Kyrgyz PGS systems.	Generalisability; cross-jurisdiction interoperability.

Phases P1–P2 are engineering-and-deployment; the central scientific phase is P3, which is when the projected outcomes of §5.2 become testable hypotheses with statistical force. This restatement is intended to make the empirical commitments of the framework explicit and evaluable.

8. Conclusion

This paper proposes a risk-triggered hybrid assurance architecture that integrates digital traceability with selective laboratory audits for organic supply chains. The contribution lies not in any individual technology component — each is well-established — but in their integration into a governance-aware compliance architecture that places AI-derived risk scores on the critical path between sensing and laboratory action, while remaining explicitly complementary to regulatory process certification.
Central Asia, and Uzbekistan in particular, offers a deployment context in which institutional newness, state-driven digitalisation and developing laboratory capacity converge to make risk-proportionate verification both necessary and feasible. The framework is presented at conceptual maturity, supported by indicative parameters from prior EU deployment and analytic projections derived from those parameters. The empirical commitments — controlled validation of accuracy, cost and consumer-confidence claims; AI-model performance reporting in field conditions; in-country feasibility studies — are stated explicitly in §6.4 and §7. We invite critique and collaboration toward refining and validating this architecture.

Funding
This research received no external funding.

Compliance with Ethical Standards
The results of this study do not involve human participants or animal subjects. No ethics-committee approval was required for this conceptual and framework-based research.

Conflicts of Interest
The authors declare no competing interests.

References

1. AgriDigital. (n.d.). *Blockchain-based commodity-backed financing pilot with Rabobank*. AgriDigital. Retrieved May 4, 2026, from https://www.agridigital.io/
2. Agroberichten Buitenland. (2025). *Agricultural trends in Uzbekistan for 2025*. Netherlands Ministry of Agriculture, Nature and Food Quality. https://www.agroberichtenbuitenland.nl/
3. Agroberichten Buitenland. (2026). *Uzbekistan's agricultural sector: Growth, exports, and investor support*. Netherlands Ministry of Agriculture, Nature and Food Quality. https://www.agroberichtenbuitenland.nl/
4. CropIn. (n.d.). *SmartFarm: AI-driven farm-management platform*. CropIn. Retrieved May 4, 2026, from https://www.cropin.com/smartfarm
5. Eshankulov, K., Zohirov, K., Bakaev, I., Shafiev, T., Nazarov, S., Temirov, Z., & Nasimov, R. (2026). An RL-enhanced multi-agent framework for scalable and intelligent business intelligence systems. *Information, 17*(3), Article 252. https://doi.org/10.3390/info17030252
6. European Commission. (2018). *Regulation (EU) 2018/848 of the European Parliament and of the Council of 30 May 2018 on organic production and labelling of organic products and repealing Council Regulation (EC) No 834/2007*. Publications Office of the European Union.
7. https://eur-lex.europa.eu/eli/reg/2018/848/oj

8. European Commission. (2024). *Regulation (EU) 2024/1781 of the European Parliament and of the Council of 13 June 2024 establishing a framework for the setting of ecodesign requirements for sustainable products (Ecodesign for Sustainable Products Regulation, ESPR)*. Publications Office of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1781/oj
9. European Commission. (2025). *Ecodesign for Sustainable Products Regulation: Work plan 2025–2030*. Publications Office of the European Union. https://commission.europa.eu/energy-climate-change-environment/standards-tools-and-labels/products-labelling-rules-and-requirements/sustainable-products/ecodesign-sustainable-products-regulation_en
10. European Union. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Publications Office of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
11. Food and Agriculture Organization of the United Nations [FAO]. (2018). *Organic agriculture in Uzbekistan: Status, practices and prospects* (Project TCP/UZB/3501). https://www.fao.org/
12. Food and Agriculture Organization of the United Nations [FAO]. (2024). *The state of food and agriculture 2024*. https://www.fao.org/publications/sofa/2024/
13. Frontiers in Sustainable Food Systems. (2024). Organic agriculture in Kyrgyzstan: Experiences with implementing participatory guarantee systems. *Frontiers in Sustainable Food Systems, 8*, Article 1453850. https://doi.org/10.3389/fsufs.2024.1453850
14. Huang, Y., Li, X., Xu, L., & Ma, Y. (2025). Digital traceability in horticulture: A systematic review of edge-cloud-blockchain-terminal (ECBT) integration with IoT and AI technologies. *Frontiers in Blockchain, 8*, Article 1636627. https://doi.org/10.3389/fbloc.2025.1636627
15. IBM. (n.d.). *IBM Food Trust: A blockchain-based food supply-chain platform*. Retrieved May 4, 2026, from https://www.ibm.com/products/supply-chain-intelligence-suite/food-trust
16. OriginTrail. (n.d.). *Decentralized knowledge graph for supply-chain traceability*. Retrieved May 4, 2026, from https://origintrail.io/
17. Shafiev, T., Bakaev, I., Eshankulov, K., & Nazarov, S. (2025). Creation of an information and analytical database based on an intelligent analysis system from online services for monitoring and forecasting the environmental condition of industrial areas. *Proceedings of SPIE, 13651*, Article 136510U. https://doi.org/10.1117/12.3061968
18. TE-FOOD International. (n.d.). *Food traceability platform*. Retrieved May 4, 2026, from https://tefoodint.com/
19. U.S. Department of Agriculture [USDA], Agricultural Marketing Service. (2023). National Organic Program (NOP); Strengthening Organic Enforcement; Final rule. *Federal Register, 88*(12), 3548–3627. https://www.federalregister.gov/d/2023-00702
20. U.S. Department of Commerce, International Trade Administration. (2025). *Uzbekistan country commercial guide: Agricultural sectors*. https://www.trade.gov/country-commercial-guides/uzbekistan-agricultural-sectors