Zero-Disclosure Material Passports for Verifiable Provenance in Multi-Tier Supply Networks

Shivani Dharmavaram
University of the Cumberlands, Kentucky, USA

***Abstract*—Global supply chains produce vast quantities of transactional data, yet most existing traceability systems force companies to choose between disclosing sensitive commercial relationships to a shared infrastructure and relying on mechanisms that do not provide strong privacy guarantees. This study introduces a zero-disclosure material passport framework for verifiable provenance in multi-tier supply networks. The framework represents product history as a Resource-Event-Agent provenance graph and authenticates each event through unlinkable aggregate signatures. Individual participant signatures collapse into a single, fixed-size aggregate proof, so the cryptographic verification object remains 192 bytes in size regardless of the supply chain depth. Selective disclosure is supported through commitment-based predicate proofs that allow regulators and auditors to verify attributes such as origin, certification status, or recycled-content thresholds without learning unrelated commercial information. A Rust prototype was evaluated on an x86-64 workstation, a Raspberry Pi 4, and a Raspberry Pi Zero 2 W. Signing latency remained below 700 ms on the most constrained device, demonstrating feasibility for low-frequency supply chain handover events and offline inspection contexts. The study further provides algorithmic descriptions of credential issuance, passport update, aggregate verification, and selective disclosure; an expanded comparison against blockchain, EDI, BBS+/anonymous-credential, and zk-SNARK-based approaches; and a discussion of the centralized Credential Authority as an explicit trust assumption. Security analysis establishes unforgeability and unlinkability under standard hardness assumptions when the Credential Authority is honest, and participant credentials are issued through the prescribed blinded protocol.**

Keywords*—*Supply chain traceability; material passports; zero-knowledge proofs; aggregate signatures; privacy-preserving authentication; circular economy; provenance verification; selective disclosure

I. Introduction

Tracking a product across a global supply chain appears straightforward in principle, but the operational reality is complex. A finished good may pass through dozens of independent companies spread across several continents, with each organization contributing value while also holding information that competitors, buyers, and even some regulators should not be able to infer. Centralized traceability databases provide a searchable record, but they also concentrate sensitive data such as supplier identities, sourcing patterns, pricing signals, and route dependencies in a single system. Permissioned blockchains distribute records across a consortium, yet their transaction ledgers often remain visible to network members. Even when payloads are encrypted, transaction metadata can expose commercially valuable relationships.

The core problem is that many traceability systems treat verifiability and disclosure as a single objective. Proving that a product history is authentic does not logically require revealing every actor involved in that history. This study separates those objectives by introducing a zero-disclosure material passport: a cryptographic digital twin of a physical product that carries a complete Resource-Event-Agent lineage while revealing only the attributes required by a particular verifier. The term zero-disclosure is used here as a system-level property. It means that the end-to-end architecture is designed to reveal nothing beyond the verifier's authorized disclosure predicate. This is broader than the cryptographic term zero-knowledge [22], which refers to a specific class of proof systems; zero-knowledge proofs are one component of the framework, but zero-disclosure is the design goal of the whole protocol.

The proposed framework combines three ideas. First, BLS aggregate signatures provide a constant-size authentication object for a chain of events. Second, re-randomizable anonymous credentials prevent an observer from linking two credential presentations to the same participant. Third, a Resource-Event-Agent provenance graph represents supply chain activity as resources, events, and role-authorized agents rather than as named companies. A verifier can therefore confirm that the right type of actor performed the right type of operation at the right point in the chain without learning that actor's identity or the surrounding commercial relationship.

The contributions of this study are as follows: First, it defines a zero-disclosure material passport architecture that integrates BLS aggregate signatures, commitment-based selective disclosure, and a Resource-Event-Agent provenance directed acyclic graph (DAG). Second, it provides algorithmic descriptions of the main protocol operations so that the method can be implemented and audited independently. Third, it states the relevant security claims and clarifies the assumptions required for unforgeability and unlinkability, including the fact that the centralized Credential Authority is trusted not to issue fraudulent or traceable credentials. Fourth, it evaluates a working prototype across three hardware platforms, from server-class infrastructure to low-power edge devices. Fifth, it expands the comparative analysis to include conventional EDI, permissioned blockchain, privacy-preserving blockchain, BBS+/anonymous-credential, and zk-SNARK-based approaches. Finally, it discusses the practical meaning of measured latency and verification costs for real industrial deployment.

The remainder of the study is organized as follows: Section II reviews related work and identifies the comparative gap. Section III presents the cryptographic framework and algorithms. Section IV describes the material passport architecture. Section V reports implementation and performance results. Section VI discusses the benchmark results, limitations,

and deployment implications. Section VII presents supply chain integration scenarios. Section VIII concludes with future research directions.

II. Related Work and Comparative Gap Analysis

Research on supply chain traceability draws from cryptographic authentication, privacy-enhancing technologies, and distributed data management. Early traceability systems relied on EDI agreements and centralized databases, which made business-to-business data exchange operationally convenient but did not provide cryptographic privacy for participants whose data were stored or processed by shared infrastructure. Digital signatures improved integrity, but conventional signatures expose the signer and are therefore poorly suited to anonymous or commercially confidential traceability [5]. In terms of the four properties used in this study, such systems may permit offline lookup against a local copy, but they provide neither participant unlinkability nor attribute-level selective disclosure, and they carry no constant-size cryptographic proof object at all.

Digital product passport research has increasingly emphasized privacy and interoperability. Structured reviews of passport systems identify privacy, traceability, lifecycle coverage, and standardization as recurring requirements that existing deployments often address only partially [7], [9], [11]. Blockchain-based provenance systems provide immutable or append-only logs, and systems such as PrivChain aim to protect provenance relationships in blockchain-enabled supply chains [4]. However, blockchain ledgers can still expose participation patterns, and many designs require online verification or proofs that grow as the provenance chain deepens. Hyperledger Fabric provides permissioned channels and strong enterprise integration properties, but channel membership and endorsement metadata can reveal participation patterns unless carefully engineered [19]. Assessed against the four target properties, these designs diverge in opposite ways: PrivChain preserves participant unlinkability and end-to-end lineage but forfeits the constant-size proof and offline verification, since its proofs grow with chain depth and its checks expect a live network, whereas Hyperledger Fabric retains workable record sizes yet surrenders unlinkability because channel and endorsement metadata expose who participated.

More advanced privacy mechanisms have also been proposed. zk-SNARKs allow a verifier to check a statement without learning witness data, and they have been applied to sustainability reporting and digital product passport verification [6]. Their main disadvantages in supply chain settings are operational complexity, circuit maintenance, trusted-setup concerns for some constructions, and limited readability for enterprise teams that must audit the workflow. BBS+ and related anonymous credential systems support selective disclosure and unlinkable presentations, making them highly relevant to privacy-preserving traceability [18]. Their challenge is that they normally address credential presentation rather than the full provenance-chain aggregation problem. A material passport needs both attribute-level disclosure and end-to-end authentication across many supply chain events. Mapped onto the four properties, both families deliver participant unlinkability and attribute-level disclosure, but neither resolves the gap on its own: zk-SNARK constructions can provide succinct verification, yet applying them to a full provenance graph shifts complexity into circuit design, setup choices for some systems, and proving cost, while BBS+ credentials address presentation rather than chain-wide aggregation and so do not supply an end-to-end DAG lineage. This is the combination that the proposed framework is designed to close, as Table I sets out.

Several provenance ontologies have been used in supply chain systems. GS1 EPCIS is widely adopted in retail and pharmaceutical tracking, but it normally records explicit business actors and read-point events. W3C PROV provides a general model for provenance in heterogeneous information systems, but it was not designed around cryptographic unlinkability. The Resource-Event-Agent ontology, originally proposed for accounting systems, models economic activity through resources, events, and agents [23]. That abstraction is useful in this framework because an agent can denote an authorized role rather than a named firm.

The gap that motivates this work is, therefore, not the absence of traceability systems, blockchains, anonymous credentials, or zero-knowledge proof systems. The gap is the absence of a single supply chain passport design that simultaneously provides: 1) a constant-size proof object, 2) offline verification, 3) participant unlinkability, and 4) attribute-level selective disclosure. The reviewed systems tend to satisfy some of these properties but not all. Table I makes this relationship explicit and connects the literature survey to the design requirements used in the proposed framework.

III. Cryptographic Framework Design

The cryptographic engine is built around BLS signatures on elliptic curve groups equipped with a bilinear pairing [5]. The important engineering property is homomorphic aggregation: signatures produced by multiple participants can be combined into a single signature whose size does not depend on the number of signers, a property also studied in efficient multi-signature designs for blockchain settings [21]. This property is essential for deep supply chains because it prevents a passport from growing linearly with the number of enterprises that have handled the product.

Let G_1 and G_2 be the pairing groups of prime order q, let G be a generator, and let $e : G_1 \times G_2 \to G_T$ be a bilinear map. Each participant i holds a secret key sk_i and public key $pk_i = sk_iG$. For an event message m, the participant computes:

$$\sigma_i = sk_i \cdot H(m) \tag{1}$$

where, $H(\cdot)$ hashes the event descriptor to the curve. The aggregate signature is:

$$\sigma = \sum_{i=1}^{N} \sigma_i \tag{2}$$

A verifier checks:

$$e(\sigma, G) = e(H(m), \sum_{i=1}^{N} pk_i) \tag{3}$$

TABLE I. COMPARISON OF THE PROPOSED FRAMEWORK AGAINST PRIOR SYSTEM FAMILIES

System family	Constant proof size	Offline verify	Participant unlinkability	Attribute disclosure	End-to-End DAG lineage	Main limitation
GS1 EDI/EPCIS [11]	No	Yes	No	No	Partial	Mature operational standard, but relies on explicit identifiers and access control.
Hyperledger Fabric [19]	Partial	No	No	Partial	Partial	Permissioned channels help governance, but metadata can reveal participation.
PrivChain [4]	No	No	Yes	Partial	Yes	Protects provenance privacy, but proof and verification costs grow with chain depth.
BBS+/anonymous credentials [16], [18]	Partial	Yes	Yes	Yes	No	Strong credential disclosure primitive, but not a complete provenance aggregation architecture.
zk-SNARK/DPP designs [6]	Partial	Partial	Yes	Yes	Partial	Expressive private verification, but circuit design, setup, and operational complexity remain high.
Proposed framework	Yes	Yes	Yes	Yes	Yes	Relies on a trusted or consortium-operated Credential Authority in the current design.

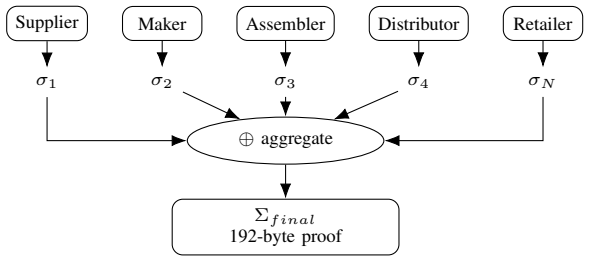

Fig. 1. Signature aggregation. Individual participant signatures are combined into one aggregate authentication proof, whose size is independent of the number of participants.

The result is a fixed-size authentication object even when N grows. In this prototype, the final Event Authentication Code is represented as a 192-byte aggregate proof. Fig. 1 illustrates how individual participant signatures collapse into this single constant-size authentication object.

Participants do not present stable identities during signing. They obtain role-specific credentials from the Credential Authority through blinded issuance and then re-randomize those credentials before each session. Re-randomization changes the public presentation while preserving the algebraic relationships needed for verification. As a result, a verifier can confirm that the signer had the required authorization without learning whether two events were signed by the same organization.

A. Protocol Algorithms

Algorithm 1 to Algorithm 3 describe the protocol in implementation-oriented terms, covering credential issuance, passport update and aggregation, and selective-disclosure verification.

Algorithm 1. Credential Onboarding and Session Randomization

Require: Participant role request R, participant key pair (sk_i, pk_i), public CA key PK_{CA}
Ensure: Re-randomized session credential $cred_i^{(t)}$
1: Participant constructs a blinded request $B(R, pk_i; r)$ using fresh randomness r.
2: Participant sends B to the Credential Authority through an authenticated enrollment channel.
3: Credential Authority checks policy eligibility for role R without learning the unblinded credential presentation.
4: Credential Authority issues a role credential $cred_i$ bound to pk_i and authorized actions.
5: For each signing session t, participant samples fresh randomness ρ_t.
6: Participant computes $cred_i^{(t)} = ReRand(cred_i, \rho_t)$.
7: **return** $cred_i^{(t)}$ for unlinkable event signing.

Algorithm 2. Passport Event Append and Aggregate Authentication

Require: Passport DAG D, event descriptor m, authorized participants $P = \{1, \ldots, N\}$
Ensure: Updated passport D' with aggregate proof Σ_{event}
1: Create an event node E_m containing timestamp, location reference, and committed operating attributes.
2: **for** each participant $i \in P$ **do**
3: Verify that $cred_i^{(t)}$ authorizes the role required by E_m.
4: Compute signature component $\sigma_i = sk_i \cdot H(m)$.
5: **end for**
6: Aggregate signature components $\Sigma_{event} = \sum_{i=1}^{N} \sigma_i$.
7: Link input resource nodes, E_m, role-authorized agent nodes, and output resource nodes in D.
8: Store Σ_{event} and inherited authentication context on the new DAG edges.
9: **return** Updated passport graph D'.

Algorithm 3. Verification and Selective Disclosure

Require: Passport D, aggregate proof Σ, disclosure predicate $P(a)$, public parameters
Ensure: Accept or reject with only authorized attributes disclosed
1: Parse the relevant resource, event, and agent nodes from D.
2: Reconstruct the event message m and authorized public-key aggregate PK_{agg}.
3: Check pairing equation $e(\Sigma, G) = e(H(m), PK_{agg})$.
4: **if** pairing check fails **then**
5: Reject the passport.
6: **end if**
7: For each requested attribute, verify the commitment opening or predicate proof.
8: Run the non-interactive Sigma-protocol verifier for $P(a)$.
9: **if** all predicates and graph links verify **then**

10: Accept and reveal only the attributes allowed by $P(a)$.
11: **else**
12: Reject.
13: **end if**

B. Selective Disclosure Construction

Each resource attribute a is stored as a commitment rather than plaintext. The prototype uses Pedersen-style commitments [17] of the form:

$$C = g^a h^r \tag{4}$$

where, g and h are public group generators and r is a uniformly sampled blinding factor. To prove a predicate $P(a)$, such as recycled content exceeding a threshold, the passport holder constructs a non-interactive Sigma-protocol proof of knowledge showing that the committed value satisfies both $C = g^a h^r$ and $P(a)$ without revealing a or r directly. The Fiat-Shamir transform makes the proof non-interactive, which is important for asynchronous supply chain settings and offline verification.

C. Threat Model and Trust Assumptions

The adversary is modeled as a probabilistic polynomial-time algorithm that may observe all public credential presentations, aggregate proofs, and protocol transcripts. It may corrupt any subset of ordinary supply chain participants and obtain their local signing keys and credentials. It may also attempt replay, double-signing, and graph-splicing attacks by recombining valid events into invalid lineages.

A central trust assumption is that the Credential Authority is not corrupted. This assumption is explicit because the Credential Authority issues the role credentials that make unlinkable authorization possible. If the Credential Authority were malicious, it could issue fraudulent credentials, refuse legitimate credentials, or embed traceable structure in credential issuance. The current design therefore provides security under an honest-but-policy-enforcing Credential Authority model. Threshold credential issuance and multi-authority issuance are treated as future work because they remove this institutional single point of trust but require additional coordination and governance.

The Credential Authority assumption does not weaken the aggregate-signature invariance property, but it does define the boundary of the privacy and authorization guarantees. In practical deployments, the Credential Authority should be operated by a regulator, standards body, industry consortium, or threshold group rather than by a single supply chain participant.

D. Security Claims

Security Claim 1 (Unforgeability): Assuming co-Diffie-Hellman hardness in the selected pairing group and proper proof-of-possession or CA-mediated key registration, an adversary that does not hold a valid credential for an authorized role cannot produce an aggregate signature component that passes the verification equation for a new event, except with negligible probability.

Reduction outline. Suppose an adversary produces a valid aggregate proof for an event requiring at least one uncorrupted authorized signer whose signature component was never issued for that message. The reduction embeds a co-CDH challenge in that target signer's public key. Signatures for corrupted or honestly simulated participants are known to the reduction and can be algebraically subtracted from the forged aggregate because aggregation is linear. What remains is a valid BLS signature on the challenge public key for a message that was not queried, yielding a solution to the underlying co-CDH/BLS forgery game. Rogue-key attacks are excluded by the CA-issued credential binding and proof-of-possession requirement during onboarding. Thus, a successful aggregate forgery implies a successful underlying BLS forgery, contradicting the hardness assumption [5].

Security Claim 2 (Unlinkability): Assuming decisional Diffie-Hellman hardness for the credential presentation group, an adversary observing two valid credential presentations cannot distinguish whether they were generated by the same participant or by two independent participants, except with negligible advantage.

Reduction outline. Each session presentation applies fresh randomization to the participant credential. The resulting group elements are distributed as randomized commitments preserving authorization relations but hiding the stable credential representation. If an adversary could link two such presentations with non-negligible advantage, a simulator could use that adversary to distinguish a valid Diffie-Hellman tuple from a random tuple by embedding the challenge tuple in the credential randomization. This contradicts the DDH assumption and follows the unlinkability logic of re-randomizable anonymous credentials [16], [18].

IV. Material Passport Architecture

A material passport is a signed, portable graph document that encodes what a verifier needs to confirm a product's lineage: what resource was handled, what event occurred, what role-authorized agent performed the action, and how that event relates to upstream and downstream resources. The graph is a DAG in which nodes and edges carry authentication context. This prevents a downstream verifier from accepting a product state unless its lineage links back through authenticated transformations. Fig. 2 summarizes the architecture connecting credential issuance, event signing, provenance-graph construction, and selective-disclosure verification.

A. Resource Nodes

A resource node captures a material, component, or product at a particular point in its lifecycle. The technical attributes it records, such as composition, grade, batch quantity, and dimensions, are held as commitments rather than as plain values. Crucially, the node keeps no clear record of the supplier, manufacturer, or handler involved. Because each attribute is committed separately, a verifier can later confirm any single claim without the holder having to reveal the complete attribute vector.

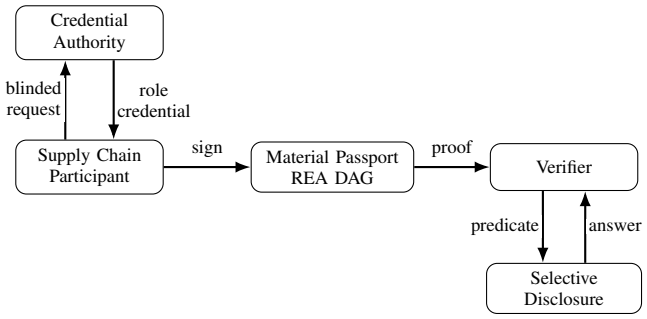

Fig. 2. Zero-disclosure material passport architecture. Blinded credential issuance, unlinkable event signing, graph-based provenance, and selective disclosure are combined in one verification workflow.

B. Event Nodes

An event node describes something that happens to a resource, whether a transformation, a measurement, an inspection, or a transfer of custody. Alongside timestamp metadata, a location reference, and the relevant operating parameters, each event carries an aggregate proof confirming that the required roles authorized it. Where an operation is complex, it can be broken down hierarchically: a sorting event, for instance, may branch by material class and then split further into sub-branches for polymer or alloy type. This layering gives auditors fine-grained insight while still allowing the passport to withhold details that a particular verifier does not need to see.

C. Agent Nodes

Agent nodes stand for the role-authorized actors in the chain. Rather than recording a company name or a tax identifier, an agent node holds an unlinkable credential commitment together with a role certificate. A verifier can therefore establish that a customs broker, recycler, certified manufacturer, or authorized inspector took part in an event, yet learn nothing about which specific organization sits behind that role.

D. Edge Authentication and Lineage Propagation

Edges in the DAG express the parent-child relationships between resources and events. When a transformation turns input resources into output resources, the resulting output node is linked both to those inputs and to the governing event node. Since the output node inherits its authentication context from the event proof, a downstream verifier can trace the chain backward without ever contacting the participants upstream. This property is what makes offline inspection workable in ports, warehouses, and field environments.

E. Physical-Digital Coupling

Each digital passport is bound to a physical item or container by a cryptographic seed embedded in a hardware identifier such as an RFID tag, a QR code, or a serialized label. Scanning that seed triggers a verification query, which runs either against a network backend or against a locally cached copy of the passport. For bulk materials, the seed can be fixed to the sealed container instead of to individual units, which preserves custody integrity at the shipment level.

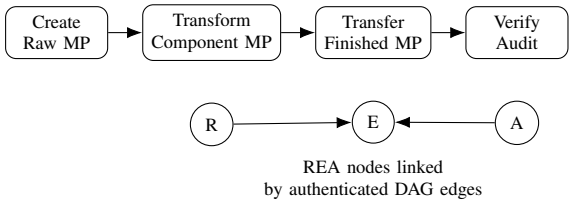

Fig. 3. Material passport lifecycle. Resource, Event, and Agent nodes are extended, as the product is created, transformed, transferred, and verified.

TABLE II. OPERATION TIMINGS AND OUTPUT SIZES

	Time (ms)			Size (B)		
Operation	X64	RPi4	RPi0	X64	RPi4	RPi0
Credential	13.7	37.3	249.2	96	96	96
Session	18.8	51.6	276.0	128	128	128
Signing	73.0	145.6	660.3	64	64	64
Verify	40.5	104.9	575.3	0	0	0
Aggregate	21.9	57.9	346.9	192	192	192
Disclosure	39.1	88.3	516.5	48–256	48–256	48–256

V. IMPLEMENTATION AND PERFORMANCE ANALYSIS

The prototype is implemented in Rust to obtain memory safety and predictable runtime behavior. Cryptographic execution is sandboxed so that signing state and credential material are isolated across operations. The provenance graph is serialized in JSON-LD, allowing semantic data and cryptographic proofs to be processed by standard graph tooling while still supporting compact verification artifacts.

System setup begins when the Credential Authority generates public parameters for the BLS12-381 curve and publishes the system verification key. Participants generate local signing keys and request role credentials through the blinded issuance protocol in Algorithm 1. At runtime, event nodes are appended using Algorithm 2; verification and selective disclosure follow Algorithm 3.

Fig. 3 summarizes the implemented passport lifecycle across creation, transformation, transfer, and verification, showing how credentials, event signatures, REA graph updates, and selective-disclosure checks are connected during execution.

The evaluation uses three hardware classes: an x86-64 workstation representing cloud or enterprise infrastructure; a Raspberry Pi 4 representing handheld or gateway devices; and a Raspberry Pi Zero 2 W representing constrained edge hardware. Table II reports operation timings and output sizes. The most resource-constrained signing operation completes in 660.3 ms, and verification completes in 575.3 ms. These values are appropriate for supply chain handovers, inspections, and batch-level traceability events, which occur at human or process timescales rather than at high-frequency sensor rates.

Table III shows scalability as participant count and stage count increase. Signing and aggregation work grow with the number of participants because each authorized role must contribute a signature component. However, the proof object carried by the passport remains fixed. The total passport size grows primarily with semantic provenance records and disclosed metadata, not with the aggregate cryptographic object.

TABLE III. SCALABILITY ANALYSIS WITH INCREASING PARTICIPANTS

Participants	Stages	Sign (ms)	Verify (ms)	Size (KB)
5	3	365	122	1.2
10	5	730	258	2.0
50	8	3,650	797	3.2
100	12	7,300	1,943	4.8
500	20	36,500	13,974	8.0
1,000	25	73,000	35,230	10.0

High-throughput distribution centers can use batch verification to amortize pairing computations across many concurrent checks. Handheld devices can use cached public parameters and pre-computed lookup tables to support offline operation. Full audit tools can recursively verify every graph edge and aggregate proof in a passport's history, which is slower but appropriate for forensic or regulatory review.

VI. DISCUSSION

The benchmark results should be interpreted in relation to the research question introduced in Section I: Can supply chain verifiability be achieved without exposing commercially sensitive participant relationships? The results indicate that the answer is yes for low-frequency provenance events such as custody transfer, batch certification, transformation completion, and inspection. The framework's central invariance property is the fixed-size aggregate proof. Even when the number of enterprises increases, the authentication object remains 192 bytes. This is the most important engineering result because it prevents the proof from becoming unusable as the supply chain deepens.

Latency matters differently across deployment tiers. On server infrastructure, signing and verification are fast enough for interactive enterprise systems. On the Raspberry Pi 4, operations remain compatible with handheld inspection and gateway-class devices. On the Raspberry Pi Zero 2 W, signing below 700 ms and verification below 600 ms are not appropriate for high-frequency sensor streams, but they are appropriate for supply chain events that occur at discrete operational checkpoints. This distinction is important: the framework is not intended to sign every raw temperature sample. Instead, it is intended to sign meaningful provenance events, such as a certified temperature excursion report or a completed cold-chain handover.

The verification cost on constrained hardware also has a deployment implication. Offline verification is feasible, but devices should pre-cache public parameters, role-policy metadata, and recent revocation information during synchronization windows. This design is appropriate for ports, warehouses, and field inspection sites where connectivity can be intermittent. When connectivity is available, online verification provides the freshest policy and revocation state. When connectivity is not available, cached verification preserves operational continuity while limiting the freshness of revocation checks.

The expanded comparison in Table I clarifies the position of the framework. BBS+ credentials and zk-SNARKs are not competitors to be dismissed; they are important privacy primitives. The proposed contribution is the integration of privacy-preserving credential presentation, fixed-size event aggregation, and graph-structured provenance into a passport architecture that supports offline verification. In this sense, the framework is an applied systems design rather than a claim that any individual primitive is new.

The most important current limitation is the centralized Credential Authority. The threat model assumes that the Credential Authority does not issue fraudulent credentials and does not embed traceable structure into credentials. This assumption is reasonable for a prototype and for some regulated settings, but it is not ideal for global industrial deployment. A malicious or compromised Credential Authority could damage both trust and privacy. Future deployments should therefore use threshold credential issuance, multi-authority governance, external audits, and transparent credential policies. Another limitation is that high-frequency sensor attestation would require batching, hardware acceleration, or a separate streaming attestation layer.

VII. SUPPLY CHAIN INTEGRATION CASE STUDIES

The framework applies to several supply chain contexts in which auditability and confidentiality are simultaneously required.

A. Pharmaceutical Supply Chains

Drug manufacturing requires evidence of ingredient origin, manufacturing conditions, packaging integrity, and handover history. A zero-disclosure passport can authenticate batch-level movement from active pharmaceutical ingredient synthesis through formulation and distribution while hiding which contract manufacturers or logistics providers participated in a particular route [3]. Regulators can verify compliance, while commercial sourcing arrangements remain private.

B. Electronics and Conflict Minerals

Electronics supply chains often span many tiers, and manufacturers may need to prove that minerals were not sourced from prohibited regions. A zero-disclosure passport can prove that each upstream tier satisfied the required certification predicate without naming every supplier. This supports conflict-mineral auditability while preserving sourcing strategy and supplier confidentiality [10].

C. Circular Economy and Material Reuse

Circular economy certification depends on proving that reused or recycled materials came from the claimed source and passed through valid transformations. The graph structure naturally supports split and merge operations: one recycled material batch may divide into multiple product streams, and multiple streams may later combine into one component. Each branch inherits cryptographic provenance from the authenticated DAG [2].

D. Food Safety and Farm-to-Fork Traceability

Farm-to-fork systems require certification of origin, pesticide use, organic status, and cold-chain integrity. Selective disclosure allows investigators to identify an affected batch during a contamination event without exposing unrelated producers or unaffected shipment data. This enables narrower recalls and reduces collateral reputational harm [15].

E. Construction Materials

Construction projects increasingly require auditable evidence of embodied carbon, recycled content, and material origin. Material passports can preserve the provenance of cement, aggregate, steel, and admixtures from source to installation. This is useful for green-building certification, circular reuse, and forensic investigation after structural failures [1].

F. Aerospace and Defense

Aerospace and defense supply chains require part-level traceability but also involve sensitive supplier relationships. A passport can prove that a component passed through authorized custody and inspection steps while concealing supplier identities when disclosure is not legally required [13]. This reduces counterfeit-part risk without unnecessarily exposing the supply network.

G. Cross-Border Trade and Customs Compliance

Customs authorities need proof of origin, export-license status, and regulatory compliance, but they do not need full commercial routing or pricing data. Selective disclosure can provide a proof package tailored to customs requirements. Because verification can operate offline, border clearance is not blocked when foreign systems are unavailable [6].

H. Precision Recall Management

Product recalls are often broader than necessary because manufacturers cannot prove which units are affected. A full provenance DAG allows a defective batch, production line, or time window to be identified cryptographically. Downstream verifiers can determine whether a unit belongs to the affected lineage without seeing unrelated production data [12].

I. Industrial IoT Integration

Industrial sensors can provide temperature, humidity, vibration, and location readings that are relevant to compliance. The framework should not sign every raw reading on constrained hardware; instead, it can sign summarized, policy-relevant events such as excursion reports, tamper alerts, or inspection confirmations. Sensor-based passport designs demonstrate the practical value of binding environmental monitoring to product records [14].

J. Regulatory Reporting Automation

Regulatory reporting often requires evidence that already exists in operational systems but is difficult to assemble and verify. A material passport can serve as the authoritative evidence graph, from which jurisdiction-specific reports are generated using selective disclosure. Each authority receives only the attestations relevant to its legal scope [8].

VIII. Conclusion and Future Directions

This study has presented a zero-disclosure material passport framework for verifiable provenance in multi-tier supply networks. The design brings together four building blocks: BLS aggregate signatures, unlinkable credential presentations, commitment-based selective disclosure, and a Resource-Event-Agent provenance DAG. Its central engineering result is a constant-size aggregate proof. However many enterprises have handled a product, the authentication object stays at 192 bytes, which is what keeps the scheme practical as supply chains grow deeper. Prototype benchmarks on an x86-64 workstation, a Raspberry Pi 4, and a Raspberry Pi Zero 2 W confirm that the approach is workable for low-frequency provenance events and for offline inspection in the field.

The framework's position relative to prior work is equally clear. Existing EDI, blockchain, BBS+/anonymous-credential, and zk-SNARK-based approaches each contribute valuable capabilities, yet none of them delivers constant-size proofs, offline verification, participant unlinkability, attribute-level disclosure, and end-to-end graph lineage at the same time. The contribution here is to combine those properties within a single passport workflow while stating the centralized Credential Authority openly as a trust assumption rather than leaving it implicit.

Several directions remain open. The most pressing is the threshold credential issuance, which would spread authority across several independent parties so that no single institution controls who may sign. Post-quantum readiness is a second priority, since pairing-based signatures will not withstand a sufficiently capable quantum adversary; lattice-based and hash-based schemes such as CRYSTALS-Dilithium [20] are natural candidates, provided they can preserve the aggregation and unlinkability the framework depends on. Beyond cryptography, further work could address high-frequency sensor attestation, the automatic extraction of structured event metadata from shipping documents, and independent security audits of open-source reference implementations. As digital product passport regulations continue to widen, zero-disclosure material passports offer a credible way to make global supply chains more accountable without forcing companies to expose the commercial relationships that keep them competitive.

References

[1] A. Vahidi, A. T. Gebremariam, F. Di Maio, K. Meister, T. Koulaeian, and P. Rem, "RFID-based material passport system in a recycled concrete circular chain," *Journal of Cleaner Production*, vol. 442, p. 140973, 2024.

[2] S. Wilson, K. Adu-Duodu, Y. Li, R. Sham, M. Almubarak, Y. Wang, E. Solaiman, C. Perera, R. Ranjan, and O. Rana, "Blockchain-enabled provenance tracking for sustainable material reuse in construction supply chains," *Future Internet*, vol. 16, no. 4, p. 135, 2024.

[3] A. Padma and M. Ramaiah, "Blockchain based solution for secure information sharing in pharma supply chain management," *Heliyon*, vol. 10, no. 22, p. e40273, Nov. 2024.

[4] S. Malik, V. Dedeoglu, S. S. Kanhere, and R. Jurdak, "PrivChain: Provenance and privacy preservation in blockchain enabled supply chains," in *Proc. 2022 IEEE Int. Conf. Blockchain*, Espoo, Finland, Aug. 2022, pp. 157–166.

[5] D. Boneh, B. Lynn, and H. Shacham, "Short signatures from the Weil pairing," *Journal of Cryptology*, vol. 17, no. 4, pp. 297–319, 2004.

[6] C. Udokwu and S. Crass, "Zero-knowledge proof extensions for digital product passports in sustainability claims reporting and verifications," *Electronics*, vol. 15, no. 4, p. 745, Feb. 2026.

[7] A. KC, S. Senaratne, S. Perera, and S. Nanayakkara, "Review of current digital technologies for material passports to enhance circularity towards net zero," *Built Environment Project and Asset Management*, ahead-of-print, 2024.

[8] D. Tortola, C. Felicioli, A. Canciani, and F. Severino, "Authenticated data visualization for hybrid blockchain-based digital product passports," *Computer Communications*, vol. 235, pp. 1–12, 2025.

[9] C. Udokwu, V. Patrick, A. Norta, and D. Draheim, "State of research on digital product passports: A structured literature review," *IEEE Access*, vol. 11, pp. 114059–114085, 2023.

[10] A. Canciani, C. Felicioli, F. Severino, and D. Tortola, "Enhancing supply chain transparency through blockchain product passports," in *Proc. 2024 IEEE Int. Conf. Pervasive Computing and Communications Workshops*, Biarritz, France, Mar. 2024, pp. 751–756.

[11] M. Jansen, T. Meisen, C. Plociennik, H. Berg, A. Pomp, and W. Windholz, "Stop guessing in the dark: Identified requirements for digital product passport systems," *Systems*, vol. 11, no. 3, p. 123, Feb. 2023.

[12] K. Berger, M. Rusch, A. Pohlmann, M. Popowicz, B. C. Geiger, H. Gursch, J.-P. Schoeggl, and R. J. Baumgartner, "Confidentiality-preserving data exchange to enable sustainable product management via digital product passports: A conceptualization," *Procedia CIRP*, vol. 116, pp. 354–359, 2023.

[13] S. De Diego and I. Gutierrez-Aguero, "Decentralized digital product passport building blocks for enhancing supply chain sovereignty and circular economy practices," *IEEE Access*, vol. 13, pp. 1–15, 2025.

[14] A. Pracucci and M. Giovanardi, "Design of a sensor-based digital product passport for low-tech manufacturing: Traceability and environmental monitoring in bio-block production," *Sensors*, vol. 25, no. 18, p. 5653, Sep. 2025.

[15] C. Lopes and J. Barata, "Digital product passport: A review and research agenda," *Procedia Computer Science*, vol. 246, pp. 981–990, 2024.

[16] J. Camenisch and A. Lysyanskaya, "An efficient system for non-transferable anonymous credentials with optional anonymity revocation," in *Advances in Cryptology – EUROCRYPT 2001*, Lecture Notes in Computer Science, vol. 2045. Berlin, Germany: Springer, 2001, pp. 93–118.

[17] T. P. Pedersen, "Non-interactive and information-theoretic secure verifiable secret sharing," in *Advances in Cryptology – CRYPTO 1991*, Lecture Notes in Computer Science, vol. 576. Berlin, Germany: Springer, 1992, pp. 129–140.

[18] S. Tessaro and C. Zhu, "Revisiting BBS signatures," in *Advances in Cryptology – EUROCRYPT 2023*, Lecture Notes in Computer Science, vol. 14008. Cham, Switzerland: Springer, 2023, pp. 691–721.

[19] E. Androulaki *et al.*, "Hyperledger Fabric: A distributed operating system for permissioned blockchains," in *Proc. 13th EuroSys Conf.*, Porto, Portugal, Apr. 2018, pp. 30:1–30:15.

[20] L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehle, "CRYSTALS-Dilithium: A lattice-based digital signature scheme," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2018, no. 1, pp. 238–268, 2018.

[21] M. Kara, A. Laouid, and M. Hammoudeh, "An efficient multi-signature scheme for blockchain," Cryptology ePrint Archive, Paper 2023/078, 2023. [Online]. Available: https://eprint.iacr.org/2023/078

[22] S. Goldwasser, S. Micali, and C. Rackoff, "The knowledge complexity of interactive proof systems," *SIAM Journal on Computing*, vol. 18, no. 1, pp. 186–208, Feb. 1989.

[23] W. E. McCarthy, "The REA accounting model: A generalized framework for accounting systems in a shared data environment," *The Accounting Review*, vol. 57, no. 3, pp. 554–578, Jul. 1982.