

A Governance-Aware, Privacy-Preserving, Event-Driven Conceptual Model for Supply Chain Traceability [†]

Aleksandar Panayotov ^{1,*} , Ivan Lambov ²  and Mariana Atanasova ^{1,3}

¹ Faculty of Mathematics and Informatics, Sofia University “St. Kliment Ohridski”, 1164 Sofia, Bulgaria; matanasova@fmi.uni-sofia.bg

² Institute of Mathematics and Informatics, Bulgarian Academy of Sciences, 1113 Sofia, Bulgaria; ilambov@math.bas.bg

³ Centre of Competence in Mechatronics and Clean Technologies “Mechatronics, Innovation, Robotics, Automation and Clean Technologies”—MIRACle, Sofia University “St. Kliment Ohridski” Laboratory, 1164 Sofia, Bulgaria

* Correspondence: alex.panayotov@fmi.uni-sofia.bg

[†] Presented at the 15th International Scientific Conference TechSys 2026—Engineering, Technologies and Systems, Plovdiv, Bulgaria, 14–16 May 2026.

Abstract

Supply chain traceability often fails in practice because relevant records are scattered across production, warehouse, transport, laboratory, and document systems. When a recall or audit is needed, firms must manually collect and reconcile evidence from many sources. Existing standards and blockchain platforms address parts of this problem, but prior work still reports recurring weaknesses in governance, confidentiality management, interoperability, and performance measurement. This paper presents a governance-aware, privacy-preserving, event-driven conceptual model for supply chain traceability. The model uses five event types—Create, Transform, Transfer, Verify, and Recall—linked through explicit lineage. It stores compact signed event headers on-ledger and anchors detailed off-ledger payloads and governance policy text through hashes. It also links data-sharing choices to consortium governance, defines validation invariants, embeds key performance indicators, and produces two regulator-ready outputs: a product passport and an audit pack. The contribution is a standards-informed conceptual artifact that integrates event semantics, provenance reconstruction, selective disclosure, governance, validation, performance measurement, and regulator-ready outputs in one cross-sector traceability model.

Keywords: supply chain traceability; conceptual model; blockchain; governance; privacy-preserving traceability; provenance; EPCIS 2.0; PROV-O



Academic Editors: Nikola G. Shakev,
Valyo N. Nikolov, Nikolay
R. Kakanakov and Sevil
A. Ahmed-Shieva

Published: 15 July 2026

Copyright: © 2026 by the authors.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Traceability systems are expected to answer a basic operational question: what happened to a product, where, when, and under whose responsibility. In practice, this is often difficult because relevant evidence is distributed across enterprise resource planning systems, warehouse and transport software, laboratory systems, spreadsheets, and exchanged documents. As a result, recalls and audits often depend on manual evidence collection and reconciliation. That process is slow, expensive, and error-prone.

Recent reviews show that distributed ledger technologies can support integrity and coordination in supply chains, but practical outcomes depend less on cryptography alone and more on standardization, governance, interoperability, and regulatory alignment [1–4].

In other words, the main problem is not simply how to store records, but how to represent events consistently, protect confidential business data, and govern change across multiple organizations. These socio-technical issues have repeatedly been identified as barriers to adoption [1–5].

Industry experience points in the same direction. The discontinuation of TradeLens showed that even a technically advanced platform can struggle if ecosystem participation and governance are not sufficiently aligned [6–8]. At the same time, policy developments such as the EU Ecodesign for Sustainable Products Regulation increase the importance of portable and verifiable product information, including Digital Product Passport-style evidence [9,10].

This paper addresses these issues by defining a governance-aware, privacy-preserving, event-driven conceptual model for supply chain traceability. The model is intentionally simple. It uses a small event vocabulary, explicit predecessor links, minimal on-ledger data, governance rules tied to operation, built-in validation checks, and regulator-ready outputs. Its purpose is to make traceability easier to reconstruct, easier to govern, and easier to evaluate. The research follows a conceptual design approach aligned with Design Science Research principles, in which a conceptual artifact is developed to address identified traceability gaps in the literature.

The contribution of the paper is not a new blockchain mechanism. It is a conceptual model that specifies what should be represented, how it should be linked, how confidentiality should be preserved, how governance should be connected to evidence exchange, and what outputs should be produced for audits and regulatory use.

2. Background, Problem, and Research Gap

GS1 EPCIS 2.0 provides a common structure for supply-chain business events such as shipping, receiving, and transformation, while W3C PROV provides a vocabulary for provenance and derivation [11,12]. Together, they are useful because operational traceability depends on two things at the same time: events must be represented consistently, and the lineage between events must remain reconstructable. These standards, however, do not by themselves define how confidential operational data should be shared across multiple organizations, how consortium governance should be tied to day-to-day traceability operation, or how traceability performance should be measured in a comparable way across implementations.

Permissioned blockchain platforms provide technical mechanisms that are relevant to those missing parts. Hyperledger Fabric supports private data collections, where only authorized parties exchange plaintext while a hash is anchored on the channel [13–15]. GoQuorum with Tessera supports private transactions [16–18]. Corda supports restricted visibility and selective disclosure to the parties involved in a transaction [19,20]. These mechanisms make it possible to prove that an event exists and has not been altered without exposing every underlying document to every participant.

Even so, prior reviews continue to identify recurring barriers. Governance remains a persistent adoption bottleneck, because traceability networks depend on agreement about membership, upgrades, responsibilities, and dispute handling. Interoperability remains difficult when participating systems use different structures, identifiers, and semantics. Performance is also hard to compare, because many reported solutions use different measures or do not define traceability indicators directly in the model [2–5,21]. The practical problem addressed in this paper is therefore broader than record storage. In real settings, traceability evidence is scattered across ERP, warehouse, transport, laboratory, and document systems; confidential information cannot be shared freely; consortium rules

are often not operationalized as part of the traceability design; and the resulting systems are difficult to compare or audit in a consistent way.

These problems explain the design choices in the proposed model. A small event vocabulary and explicit predecessor links address fragmented records and ambiguous provenance [11,12]. Selective on-/off-ledger data placement addresses confidentiality and data minimization [13,14,16–20]. A governance kit addresses membership, upgrades, disputes, emergencies, and audit access [2–5]. Validation invariants address broken or contradictory trace states. KPI instrumentation addresses weak comparability [2–4]. Product passports and audit packs address the need for portable evidence that can be used by auditors and regulators [9,10].

The research gap follows from this combination of requirements. Existing work addresses important subsets of the problem, but the literature does not appear to provide a single publicly documented conceptual artifact that combines all of the following in one model: a compact event grammar, explicit lineage reconstruction, privacy-preserving selective disclosure, governance integrated into the operating model, embedded validation and KPI logic, and dedicated regulator-ready outputs. The contribution of this paper is to define such an artifact.

To clarify this contribution, Table 1 compares the proposed model with three close traceability models. The table focuses only on dimensions that are central to the claimed contribution of this paper.

Table 1. Comparison with close traceability models in the literature.

Work	Explicit Event/Business-Event Model	Explicit Lineage/History Mechanism	Privacy/Selective Disclosure as a First-Class Design Element	Explicit Governance Concept Integrated into the Model	Embedded KPI/Indicator Logic	Dedicated Regulator-Facing Output Artifacts
Li et al. [22]	✓	✗	✓	✗ *	✗	✗
Dietrich et al. [23]	✓	✓	✗	✓	✗	✗
Li et al. (ProChain) [24]	✗ **	✗	✓	✗	✗	✗
This paper	✓	✓	✓	✓	✓	✓

* Li et al. (2022) include permission and membership management through the underlying blockchain platform, but not a governance concept integrated into the traceability model at the same level as the governance kit proposed here. ** Li et al. (2024, ProChain) model participant stages and traceability flow, but their core contribution is privacy protection through pseudonyms, encryption, and smart-contract-based access control rather than a reusable event grammar or business-event model.

The comparison shows that close prior works address important subsets of the problem but not the full combination targeted here. Li et al. [22] combine EPCIS/CBV-based event handling with a hybrid on-chain/off-chain privacy architecture, which makes that work close on interoperability and confidentiality. Dietrich et al. [23] provide a clear event model, explicit history links, and an integrated governance concept, which makes that work close on operational structure and lineage. ProChain [24] is close to privacy-preserving traceability because it treats selective disclosure and protected access as first-class concerns. However, none of the compared works combines these aspects with embedded validation and KPI logic and with dedicated outputs such as a product passport and audit pack. The contribution of the present paper therefore lies not in any single element, but in integrating these elements into one compact, governance-aware, privacy-preserving, event-driven conceptual model.

This framing is important for the scope of the paper. The contribution is not a full implementation claim nor a new blockchain platform. It is a conceptual model that specifies what should be represented, how it should be linked, how confidentiality should be preserved, how governance should be connected to evidence exchange, how trace states should be validated, how performance should be measured, and what outputs should be

produced for operational and regulatory use. The value of the model lies in making these interdependent requirements explicit in one design artifact.

3. Conceptual Model

3.1. Purpose, Scope, and Actors

The model is a governance-aware, privacy-preserving, event-driven traceability framework for multi-party supply chains. It covers five actor groups: producers, processors, transport or storage providers, retailers or distributors, and public authorities. It tracks three main object types: lots or batches, shipments, and attestations such as certificates or laboratory results. Its purpose is to represent the life cycle of a lot or shipment in a form that is traceable, auditable, and compatible with selective disclosure.

3.2. Inputs, Process, and Outputs

The main inputs are event submissions from enterprise systems such as ERP, warehouse, transport, and laboratory software; organizational identities and roles; governance policies; and optional telemetry or attestation data. Each submission contains a compact signed header and an off-ledger payload linked by its hash. The header carries the minimum information needed for ordering, integrity, and lineage. The payload contains the detailed underlying record.

The model process is straightforward. Events are submitted, validated, linked to predecessor events, and added to the lineage graph. Governance rules determine who may submit, read, or change model rules. Invariants prevent invalid states. The resulting graph supports trace queries and produces practical outputs, as summarized in Figure 1.




Figure 1. Overview of the governance-aware, privacy-preserving, event-driven traceability model, showing inputs, core components, internal process, and outputs.

The outputs are a *product passport*, an *audit pack*, and a *set of KPIs* computed directly from the event graph. The product passport summarizes the status and evidence state of a lot or shipment [9,10]. The audit pack provides ordered headers, proofs, and references to detailed records. The KPIs provide a structured basis for evaluating traceability performance [2–4].

3.3. Event Vocabulary and Lineage

The model uses five event types. *Create* brings a lot into scope. *Transform* records that one or more input lots become one or more output lots. *Transfer* records custody or location change. *Verify* records an attestation such as a laboratory result or certificate reference. *Recall* marks an item as under investigation and blocks further transfers until the issue is cleared.

Each event records who performed it, when it happened, what item it concerns, and which earlier events led to it. These links form a directed acyclic graph, as illustrated in Figure 2. This lineage structure allows backward tracing to origin and forward tracing to affected descendants. The event vocabulary is intentionally small because the goal is to cover essential supply chain operations while remaining understandable and reusable across settings. The mapping to EPCIS and PROV preserves semantic interoperability: Create and Transfer correspond to EPCIS business events, Transform corresponds to a TransformationEvent, and lineage links correspond to PROV derivation relations [11,12].




Figure 2. Example lineage graph showing Create, Transform, Transfer, Verify, and Recall events linked by predecessor relationships.

3.4. Data Placement and Selective Disclosure

The model separates compact event headers from full payloads. Headers are stored on the ledger because they provide ordering, signatures, and integrity proof. Full payloads remain off-ledger because they may contain confidential operational or commercial details. The link between the two is the payload hash. This design supports verifiable evidence without forcing full disclosure [13,14,16–20,25], as illustrated in Figure 3.

A typical header contains an event identifier, event verb, asset identifier, timestamp, actor identifier, predecessor links, payload hash, payload schema identifier, and actor signature. Typical off-ledger payloads include production orders, transport documents, laboratory results, certificates, and recall authorizations.




Figure 3. On-/off-ledger data placement and selective disclosure in the traceability model.

The basic logic is simple: event headers and content hashes are recorded on the ledger; full document payloads remain off-ledger; governance policy text remains off-ledger but its hash is anchored on the ledger. This arrangement preserves confidentiality and data minimization while keeping proof of existence, integrity, and policy state [13,14,16–20].

3.5. Governance Kit

Governance is included because traceability depends not only on event recording, but also on who may participate, who may change schemas or rules, how disputes are resolved, and how exceptional cases are handled. The model therefore treats consortium policy as part of the technical design. Policy text is versioned off-ledger, and its hash is anchored on-ledger so that participants can verify which rule set was active when a given event was recorded.

The governance kit covers five domains: *membership*, *change control*, *disputes*, *emergencies*, and *audit access*. Membership rules define how organizations enter or leave the consortium. Change-control rules govern how schemas and contracts evolve. Dispute rules specify evidence requirements and response timelines. Emergency rules govern urgent corrective actions. Audit-access rules define what regulators and auditors may inspect. This makes governance visible, versioned, and evidentially linked to operational traceability rather than informal.

3.6. Validation Invariants

The model contains four basic checks. **Custody continuity** means that only the current custodian may record a transfer. **Quantity conservation** means that a transformation must balance inputs and outputs within tolerance. **Verification integrity** means that a Verify event must reference a valid attestation payload hash. **Recall lock** means that an item under recall cannot be transferred until cleared.

These checks reduce the chance that broken or contradictory records remain undetected until an audit. Their purpose is to move error detection earlier in the traceability process, which should reduce audit effort and improve trace reliability.

3.7. KPI Instrumentation

The model embeds three indicators directly in its data logic. **Time-to-trace** is the time between the queried event and the earliest Create event on the lineage path. **Audit**

hand-offs count organization changes along the minimal evidence path. *Dispute cycle time* measures the time between dispute opening and resolution. Defining these measures in the model matters because prior reviews note that traceability projects often use inconsistent performance metrics, which makes comparison difficult [2–4,26].

These indicators were chosen as the minimum useful set for a cross-sector conceptual core because they measure the three performance dimensions that matter in almost any traceability setting: how quickly a trace can be reconstructed, how much effort is needed to assemble auditable evidence, and how long it takes to resolve contested records.

3.8. Outputs and Practical Use

The model produces two main evidence artifacts. The *product passport* is a compact signed view of a lot or shipment, including identifiers, current status, recent location, key attestations, and any open recall. The *audit pack* is a verifiable package containing the scope of the request, the ordered list of event headers, signatures or equivalent proofs, and references to off-ledger records. These outputs are aligned with the broader policy direction toward portable product information and regulator-ready evidence [9,10].

A simple end-to-end example shows how the model works. A retailer receives a report of possible contamination for item X. A laboratory issues a Verify event referencing the analysis. The retailer opens a Recall for the affected lots. Because the model uses explicit lineage, the system can identify descendant lots linked to the original Create event. Because recall lock is active, further transfers are blocked. An auditor can then run a trace to origin and generate an audit pack for the regulator, who can verify signatures and hashes without requesting every internal document. Table 2 illustrates this end-to-end scenario step by step, showing the on-ledger header, off-ledger payload, applicable invariant, and resulting state at each step.

Table 2. Step-by-step walkthrough of the contamination scenario, consistent with the lineage graph in Figure 2.

Step	Event	Actor	On-Ledger Header	Off-Ledger Payload	Invariant Check	Resulting State
1	Create C1	Producer	Header for Lot A: event ID, Create verb, asset ID, timestamp, actor signature, payload hash, policy hash	Production record	—(initial creation)	Lot A registered
2	Transform T1	Processor	Header for Lot C with predecessors = {C1, C2}	Production order linking input Lots A and B to output Lot C	Quantity conservation	Lot C derived from Lots A and B
3	Transfer TR1	Carrier	Header for shipment containing Lot C, predecessor = {T1}	Transport document	Custody continuity	Custody changes from processor to carrier
4	Verify V1	Laboratory	Header referencing the shipment, predecessor = {TR1}	Laboratory result attestation	Verification integrity	Attestation linked to shipment
5	Recall R1	Retailer	Header marking shipment under recall, predecessor = {V1}	Recall authorization	Recall lock activated	Further transfers blocked

3.9. Assumptions and Constraints

The model operates under several assumptions. Participants maintain synchronized clocks, preserve off-ledger payloads, harmonize identifiers, manage signing keys securely, and follow consortium decisions. It also assumes that digital artifacts such as signed passports and audit packs are acceptable for audit or regulatory use.

The model also has clear constraints. Only headers and hashes are placed on-ledger, which improves confidentiality but limits on-chain detail. The schema is intentionally small, which improves reuse but reduces sector-specific richness unless extensions are approved. The design assumes permissioned blockchain infrastructure with privacy and access control. Finally, some governance outcomes still depend on participant compliance and off-chain agreements.

The model should therefore be understood as a cross-sector conceptual core that requires sector-specific schema extensions, governance tailoring, and empirical validation before operational deployment.

4. Discussion and Contribution

The model addresses the traceability problem through a direct causal chain. A small standardized event vocabulary reduces heterogeneity in record keeping, and explicit predecessor links reduce ambiguity in provenance reconstruction [11,12]. Minimal on-ledger data combined with hashed off-ledger payloads supports proof without unnecessary exposure [13–19]. Governance policies tied to technical operation reduce ambiguity around participation, upgrades, disputes, and emergencies, while KPI instrumentation improves comparability [2–5]. Product passports and audit packs convert internal trace data into portable outputs for operational and regulatory use [9,10].

Industrial deployment experience reinforces this design rationale. Industrial platforms such as IBM Food Trust and the now-discontinued TradeLens illustrate the practical importance of governance alignment and ecosystem participation alongside ledger technology [6–8]—concerns that the proposed model elevates from operational background to integrated design element.

The contributions of the paper are therefore:

- Governance-aware, privacy-preserving, event-driven conceptual model for supply chain traceability.
- Explicit lineage logic for reconstructable provenance across organizations.
- Standards-informed alignment with EPCIS/CBV and PROV.
- Governance integrated directly into evidence exchange.
- Privacy-preserving on-/off-ledger evidence architecture.
- Validation, KPI logic, and regulator-oriented outputs: product passport and audit pack.

Unlike approaches in which governance remains implicit or external, the present model links governance directly to evidence exchange by anchoring versioned policy state to the ledger and by defining operational rules for membership, change control, disputes, emergencies, and audit access.

As a conceptual contribution, the model is assessed in this paper primarily through analytical adequacy rather than empirical deployment. Its value lies in whether the proposed elements jointly address the identified traceability problem: consistent event representation, explicit lineage reconstruction, confidentiality-preserving evidence exchange, governance linked to operation, validation of trace states, comparable performance indicators, and regulator-ready outputs. To make this analytical adequacy explicit, Table 3 maps the core traceability requirements from the problem statement to the model elements that address them. The paper therefore claims conceptual coherence, standards-informed interoperability, and practical plausibility, but does not claim full standards conformance, prototype

validation, or sector-specific implementation. These remain appropriate directions for future evaluation.

Table 3. Traceability requirements and their realization in the proposed model.

Traceability Requirement	Model Element(s)	How the Requirement Is Addressed
Consistent representation of essential supply-chain events	Five-event vocabulary	The model uses a compact event grammar— <i>Create, Transform, Transfer, Verify, Recall</i> —to represent the core operational events needed for traceability across organizational boundaries.
Reconstructable provenance across organizations	Predecessor links; lineage graph	Each event may reference prior events, forming an explicit lineage structure that supports backward tracing to origin and forward tracing to affected descendants.
Confidentiality-preserving evidence exchange	Minimal on-ledger headers; hashed off-ledger payloads	The model separates integrity proof from full disclosure by storing only compact signed headers on-ledger, while preserving detailed evidence off-ledger under hash linkage.
Governable multi-party operation	Governance kit; policy versioning and hash anchoring	Membership, change control, disputes, emergencies, and audit access are defined explicitly, and the active policy state can be verified through anchored policy hashes.
Prevention of invalid trace states	Validation invariants	The model includes invariants such as custody continuity, quantity consistency, verification integrity, and recall lock to detect broken or contradictory trace states.
Comparable traceability performance evaluation	KPI instrumentation	The model embeds a minimum cross-sector indicator set—such as time-to-trace, audit hand-offs, and dispute cycle time—to support structured comparison across implementations.
Regulator- and auditor-usable outputs	Product passport; audit pack	Internal trace records are converted into portable outputs that can support operational review, audit preparation, and regulatory inspection.
Semantic interoperability with existing standards	EPCIS/CBV-aligned event semantics; PROV-aligned lineage logic	The model is conceptually grounded in recognized standards for event representation and provenance, while remaining a conceptual artifact rather than a full conformance implementation.

Qualitative expert evaluation is planned as future work to assess whether the event grammar, governance kit, privacy design, and outputs are regarded as adequate and practically useful by experts from production, logistics, retail, compliance, and ledger architecture. Future work also includes a more formal specification of the governance model—for example, expressing the policy lifecycle and governance decisions as a state model—and a prototype implementation to validate the architectural claims empirically.

5. Conclusions

This paper presented a governance-aware, privacy-preserving, event-driven conceptual model for supply chain traceability. It addresses a recurring practical problem:

traceability evidence is often fragmented across organizational and technical boundaries, while governance is difficult and confidential data cannot be shared indiscriminately.

The proposed model responds to this problem through a compact but integrated design. It combines five event types, explicit lineage, minimal on-ledger headers with hashed off-ledger payloads, a governance kit, validation invariants, KPI instrumentation, and two practical outputs: a product passport and an audit pack. In this way, the model is intended to make traceability easier to reconstruct, easier to govern, easier to verify, and easier to compare across implementations.

The main contribution of the paper lies not in a new blockchain platform or consensus mechanism, but in defining a conceptual artifact that brings together functions that are often treated separately: event representation, lineage reconstruction, confidentiality-preserving evidence exchange, governance, validation, performance measurement, and regulator-ready outputs. The paper therefore contributes a cross-sector design core for traceability systems that must operate under both interoperability and confidentiality constraints.

The model is not presented as a full implementation or conformance claim. Rather, it should be understood as a standards-informed conceptual basis that can guide sector-specific extensions, governance tailoring, and future empirical evaluation. The next step is qualitative expert assessment to determine whether the model's event grammar, governance kit, privacy design, and outputs are regarded as adequate and practically useful in real multi-party settings.

Author Contributions: Conceptualization, A.P.; methodology, A.P. and M.A.; investigation, A.P.; writing—original draft preparation, A.P.; writing—review and editing, M.A. and I.L.; supervision, M.A.; project administration, M.A. and I.L.; funding acquisition, M.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Bulgarian Ministry of Education and Science under the National Program “Young Scientists and Postdoctoral Students—2”. The authors gratefully acknowledge the support provided by the project UNITE BG16RFPR002-1.014-0004 funded by PRIDST. This work was supported by the Center of Competence for Mechatronics and Clean Technologies “Mechatronics, Innovation, Robotics, Automation and Clean Technologies”—MIRACle, with the financial support of contract No. BG16RFPR002-1.014-0019-C01, funded by the European Regional Development Fund (ERDF) through the Programme “Research, Innovation and Digitalisation for Smart Transformation” (PRIDST) 2021-2027.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: The help of artificial intelligence was used in the creation of this paper.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Saberi, S.; Kouhizadeh, M.; Sarkis, J.; Shen, L. Blockchain technology and its relationships to sustainable supply chain management. *Int. J. Prod. Res.* **2018**, *57*, 2117–2135. [\[CrossRef\]](#)
2. Han, Y.; Fang, X. Systematic review of adopting blockchain in supply chain management: Bibliometric analysis and theme discussion. *Int. J. Prod. Res.* **2023**, *62*, 991–1016. [\[CrossRef\]](#)
3. Hübschke, M.; Buss, E.; Holschbach, E.; Lier, S. Blockchain in supply chain management: A comprehensive review of success measurement methods. *Manag. Rev. Q.* **2025**, 1–55. [\[CrossRef\]](#)
4. Karaduman, Ö.; Gülhas, G. Blockchain-Enabled Supply Chain Management: A Review of Security, Traceability, and Data Integrity Amid the Evolving Systemic Demand. *Appl. Sci.* **2025**, *15*, 5168. [\[CrossRef\]](#)

5. Bernards, N.; Campbell-Verduyn, M.; Rodima-Taylor, D. The veil of transparency: Blockchain and sustainability governance in global supply chains. *Environ. Plan. C Politi-Space* **2022**, *42*, 742–760. [CrossRef]
6. A.P. Moller-Maersk. A.P. Moller—Maersk and IBM to Discontinue TradeLens, a Blockchain-Enabled Global Trade Platform. 29 November 2022. Available online: <https://www.maersk.com/news/articles/2022/11/29/maersk-and-ibm-to-discontinue-tradelens> (accessed on 11 July 2026).
7. IBM Support. TradeLens Discontinued. 31 March 2023. Available online: <https://www.ibm.com/mysupport/s/topic/0TO50000000IQPpGAO/tradelens> (accessed on 11 July 2026).
8. Reuters. Maersk, IBM Discontinue Shipping Blockchain Platform. 29 November 2022. Available online: <https://www.reuters.com/technology/maersk-ibm-discontinue-shipping-blockchain-platform-2022-11-29/> (accessed on 11 July 2026).
9. 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, Amending Directive (EU) 2020/1828 and Regulation (EU) 2023/1542 and Repealing Directive 2009/125/EC (Text with EEA Relevance). Available online: <https://eur-lex.europa.eu/eli/reg/2024/1781/oj> (accessed on 11 July 2026).
10. European Commission. Ecodesign for Sustainable Products Regulation. 2024. Available online: https://commission.europa.eu/energy-climate-change-environment/standards-tools-and-labels/products-labelling-rules-and-requirements/ecodesign-sustainable-products-regulation_en (accessed on 11 July 2026).
11. GS1. EPCIS and CBV Implementation Guideline. Release 2.0, Ratified Mar. 2023. Available online: <https://ref.gs1.org/guidelines/epcis-cbv/> (accessed on 11 July 2026).
12. W3C. PROV-O: The PROV Ontology. W3C Recommendation. 30 April 2013. Available online: <https://www.w3.org/TR/prov-o/> (accessed on 11 July 2026).
13. Hyperledger Fabric Documentation. Private Data. Available online: <https://hyperledger-fabric.readthedocs.io/en/latest/private-data/private-data.html> (accessed on 11 July 2026).
14. Hyperledger Fabric Documentation. Private Data Architecture Details. Available online: <https://hyperledger-fabric.readthedocs.io/en/latest/private-data-arch.html> (accessed on 11 July 2026).
15. Androulaki, E.; Barger, A.; Bortnikov, V.; Cachin, C.; Christidis, K.; De Caro, A.; Enyeart, D.; Ferris, C.; Laventman, G.; Manevich, Y.; et al. Hyperledger fabric: A distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference, Porto, Portugal, 23–26 April 2018*; ACM: New York, NY, USA, 2018; pp. 1–15. [CrossRef]
16. ConsenSys GoQuorum Private Transactions. 2023. Available online: <https://goquorum.readthedocs.io/Privacy/Lifecycle-of-a-private-transaction/> (accessed on 11 July 2026).
17. ConsenSys GoQuorum Tessera. Available online: <https://goquorum.readthedocs.io/Privacy/Tessera/Tessera/> (accessed on 11 July 2026).
18. ConsenSys GoQuorum Transaction and Contract Privacy. Available online: <https://goquorum.readthedocs.io/Privacy/Overview/> (accessed on 11 July 2026).
19. R3 Corda Documentation. Identity: Party and AnonymousParty. Version 4.12/5.1. Available online: <https://docs.r3.com/en/platform/corda/4.12/community/api-identity.html> (accessed on 11 July 2026).
20. R3 Corda Documentation. Notaries. Version 5.1. Available online: <https://docs.r3.com/en/platform/corda/5.1/developing-applications/ledger/notaries.html> (accessed on 11 July 2026).
21. Dasaklis, T.K.; Voutsinas, T.G.; Tsoulfas, G.T.; Casino, F. A Systematic Literature Review of Blockchain-Enabled Supply Chain Traceability Implementations. *Sustainability* **2022**, *14*, 2439. [CrossRef]
22. Li, L.; Qu, H.; Wang, H.; Wang, J.; Wang, B.; Wang, W.; Xu, J.; Wang, Z. A Blockchain-Based Product Traceability System with Off-Chain EPCIS and IoT Device Authentication. *Sensors* **2022**, *22*, 8680. [CrossRef] [PubMed]
23. Dietrich, F.; Louw, L.; Palm, D. Blockchain-Based Traceability Architecture for Mapping Object-Related Supply Chain Events. *Sensors* **2023**, *23*, 1410. [CrossRef] [PubMed]
24. Li, J.; Wang, Z.; Guan, S.; Cao, Y. ProChain: A privacy-preserving blockchain-based supply chain traceability system model. *Comput. Ind. Eng.* **2023**, *187*, 109831. [CrossRef]
25. Sezer, B.B.; Topal, S.; Nuriyev, U. TPPSUPPLY: A traceable and privacy-preserving blockchain system architecture for the supply chain. *J. Inf. Secur. Appl.* **2022**, *66*, 103116. [CrossRef]
26. Pang, S.; Teng, S.W.; Murshed, M.; Van Bui, C.; Karmakar, P.; Li, Y.; Lin, H. A survey on evaluation of blockchain-based agricultural traceability. *Comput. Electron. Agric.* **2024**, *227*, 109548. [CrossRef]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.