Traceability in Modern Manufacturing

Devox Research Paper

Author(s): Konstantyn Gytko, CEO
Affiliation: Devox Software (DXSoft, Inc.)
Tags: AI & ML, Manufacturing, Trends
Document type: Industry whitepaper (not peer-reviewed)
Date: August 2026
DOI: *[Assigned automatically by ResearchGate upon upload]*
Supplementary materials: https://drive.google.com/drive/folders/1ydH7wj4-P_SaEFiGcRXHUonnrDF-bEcX

Summary. Modern manufacturing traceability turns fragmented production records into a unified, auditable data foundation that supports end-to-end genealogy, regulatory compliance, and trustworthy AI-driven decision-making.

Abstract

In the past, manufacturing traceability was designed to track recalls, genealogy, and quality control for materials, parts, and ready-to-ship products through every stage of production and the supply chain. Now, thanks to industry development, it is evolving into a broader information infrastructure. As batch and lot identification fail to keep pace with this new reality, manufacturers are leaning toward persistent product identities, event histories, and more.

This paper proposes a reference architecture for manufacturing traceability. It's based on data capture, integration and normalization, analytics and AI, and compliance according to the European Union's Ecodesign for Sustainable Products Regulation (ESPR) standards. This model helps to distinguish digitization effort from integrated traceability results.

Keywords

manufacturing traceability, digital product passport, MES, ERP, Industry 4.0, AI readiness, supply-chain traceability, product genealogy, industrial AI, manufacturing compliance

1. Introduction

1.1 Background

In its core, traceability answers narrow manufacturing questions: “Which batch produced this item?” or “Which materials entered that batch?” Disconnected, these questions offer a limited perspective. Modern manufacturing environments demand more than the genealogy of a finished product but also the material origin, suppliers' portfolio, real production conditions and histories, and so on, all across the product lifecycle. This transition is fueled by tech as much as the business reality.

Firstly, the regulating norms expect the manufacturing and product history to be available in structured and verifiable forms. In particular, the EU Ecodesign for Sustainable Products Regulation establishes the framework for Digital Product Passports, where it states where and how the information must be contained at the model, batch, or item level; who may access it, and how long it must remain available.

Secondly, in 2026 we expect the definitive phase of the EU Carbon Border Adjustment Mechanism to begin, according to which emissions of commercial and operational information become viable for compliance.

Thirdly, some sector-specific traceability requirements continue to mature. For example, in the automotive industry, manufacturers are expected to link incoming component lots, suppliers, production orders, processes, inspections, and shipments according to the IATF 16949 and OEM-specific requirements.

Fourthly, AI enhances traceability data value. As traditional systems retrieve records from the system after an event, AI-enabled software can potentially identify anomalies and deviations and assist specialists in real time.

These trends are fueled by manufacturing data. Identifiable, contextualized, interoperable, historically reconstructable, and trustworthy production data is in the middle of it.

1.2 Problem Statement

Isolated operational objectives are no longer relevant. In 2026, a finished batch can be associated with a production order; however, the complete chain reconstruction is impossible. You can't tell the origin of the material, storage conditions, etc.

Often, it becomes a problem, more pronounced when it goes to sophisticated supply chains and retail networks. For instance, to investigate who is liable for the materials of inferior quality, you need to check all the supply chain links. The other challenges may occur in situations when supplier identifiers don't correspond with internal material identifiers or MES timestamps differ from historian timestamps, or quality evidence exists only in PDFs and spreadsheets.

Two major risks are associated with this fragmentation: a compliance fragmentation when the evidence exists but can't be efficiently reported and an AI fragmentation when manufacturers can't reliably establish the context despite large volumes of manufacturing data in control.

For this reason, the central research problem examined in this paper is how to structure traceability architecture for the trusted data foundation to support operational genealogy and AI-driven analysis.

1.3 Contributions

This paper makes three practical contributions in traceability in modern manufacturing:

- a 5-layer reference architecture across shop-floor capture, integration, traceability records, AI, and reporting
- a standards and regulatory mapping framework
- a 5-level AI readiness maturity model

The architecture is deliberately technology-agnostic. We don't recommend a specific cloud provider and/or AI model.

2. Background Notions: The Current State as of 2026

2.1 Traceability in Manufacturing: Now and Back Then

Manufacturing traceability has evolved through several overlapping technological generations. The following table describes and compares the increments of the manufacturing traceability evolution over the years.

Dimension	Traditional Traceability	Modern Traceability	Examples
Traceability Units	Batch or lot	Batch, lot, serialized item, component, shipment, and individual production event	Then: 5,000 bearings traced to Lot B-204. Now: Bearing SN-18472 is linked to its supplier lot, assembly, machine, process parameters, inspection results, and finished vehicle.
Main Records	Papers, batch sheets, inspection forms, physical labels	MES/ERP records, IoT events, digital certificates	Then: An operator signs a paper/digital inspection sheet. Now: MES records the inspection automatically and links the result, timestamp, and measurements.
Production Genealogy	Often reconstructed manually after an incident	Built continuously from material, process, equipment, inspection, rework, and shipment events through scanners, PLCs, sensors, RFID, machine vision, etc.	Then: Engineers go into production logs to identify products that may contain a defective component. Now: A genealogy query shows all affected serial numbers with their production histories.
Cross-System Integration	Limited	APIs, middleware, event streaming	Then: Supplier data is in ERP, while machine data is in SCADA and inspection results in QMS. Now: ERP supplier lot is linked with MES work order and PLC process data backed by QMS inspection.
Supply-Chain Visibility	Inside one plant	Extends across suppliers, manufacturers, distributors, and customers	Then: The manufacturer knows when a material entered its plant and which batch used it. Now: Traceability connects supplier origin with incoming material and its transformation into a finished product and shipment to a customer.
Standards	Internal conventions	General ISA-95, OPC UA, GS1 standards, and sector-specific data standards	Then: “Part 348A” is readable inside Factory A only. Now: Standardized identifiers and event structures let exchanging product and process information between sites and systems.
Analytics	Historical reports and basic KPIs	Predictive quality, anomaly detection, graph analytics, computer vision, and root-cause analysis	Then: A monthly report shows, for instance, a 2.1% defect rate. Now: Analytics can reveal that defects correlate within the link “Supplier Lot X +

Dimension	Traditional Traceability	Modern Traceability	Examples
			Machine 7 + a specific process condition."
AI Use	Little or none	Positive	Then: Engineers manually review production records. Now: AI analyzes inspection histories to rank likely causes to support human judgment.

Table 1. The comparison of traditional and modern traceability features.

This comparison mirrors the shift from static documentation to context-specific histories. Earlier methods depended primarily on paper batch records, manually maintained inventory ledgers, and physical labels; however, a batch record often proves that a component could have been used in a product but lacks the complete event history.

RFID, IoT, and serialized identity are other methods that appeared to increase the volume and granularity of manufacturing events. With them, instead of recording only that a batch, manufacturers could potentially record the sequence.

Furthermore, a modern product is not merely associated with attributes. It moves, transforms, passes from one handler to another, becoming a problem for representing supply-chain and product events from a perspective of both manufacturers and regulators. In particular, Electronic Product Code Information Services (EPCIS) directly separates event data from master data.

Another Way: Distributed-Ledger Technology (DLT)

With the rise of blockchain and distributed-ledger architectures, manufacturers introduced yet another model: tamper-evident records can be shared securely among supply-chain participants without any central authority. Although DLT can't clean inaccurate source data, it has proved its value in multi-party environments. Yet it remains an optional tool rather than a mandatory foundation.

2.2 Regulations and Compliance Outlook

Regulations differ from state to state; however, there are global tendencies with overlapping requirements for product identity, sustainability, safety, and supply-chain due diligence as follows:

EU Digital Product Passport (DPP). Regulation (EU) 2024/1781, the Ecodesign for Sustainable Products Regulation (ESPR), establishes the DPP framework, according to which DPP information must be suitable for interoperable environments with, for instance, persistent unique product identification. It also specifies open standards to passport information.

Moreover, ESPR doesn't apply uniformly to every manufactured product, as the regulation only establishes the legal and technical framework for DPP, while the detailed requirements for individual product groups are contained in the subsequent delegated acts.

Consequently, manufacturers should distinguish between the ESPR/DPP framework and the exact data requirements for a particular product group. Particularly, a manufacturer of batteries and a manufacturer of textiles face different timelines and disclosure requirements. From an architecture perspective, flexibility gains even greater importance here as manufacturers need to trace data across existing products, suppliers, processes, and so on instead of just complying with one standard.

Carbon Border Adjustment Mechanism (CBAM). From January 1, 2026, for manufacturers and importers dealing with covered goods, traceability therefore intersects with emissions accounting. The outstanding challenge is not storing a carbon figure but connecting emissions to operations and facilities (as well as production methods, materials, etc.)

Supply-chain due diligence. Supply chain due diligence frameworks add liability to track who supplied material and under what circumstances. For instance, the U.S. Uyghur Forced Labor Prevention Act highlights supply-chain mapping and documentation for importers as standard practice. Another act, the EU Corporate Sustainability Due Diligence Directive, is similarly a step toward formalized value-chain due diligence.

Automotive Traceability. As one of the specific industry examples, IATF 16949 environments demand connecting safety-critical components and production information across complex supplier networks, while OEM-specific requirements add extra obligations.

In all cases, an architectural framework can't be underestimated, uniting all requirements into a foundation for safe and compliant manufacturing operations.

2.3 Traceability Architectures Variations

In our view, 4 main architectural patterns dominate across modern organizations. They differ primarily in where the authoritative traceability records reside and how events are captured and stored across systems and for further use in analytics and AI.

Centralized ERP/MES Architecture

In a centralized architecture, ERP or MES acts as the principal repository for traceability information. ERP typically provides the business context about the materials, suppliers, and so on, while MES layers up shop-floor information with operators and quality results.

Aspect	**Centralized ERP/MES Architecture**
Strengths	Simple governance Based upon existing material and production data Good transaction controls Easy connections Fewer platforms to govern
Weaknesses	Difficult to incorporate high-frequency sensor data Often dependent on proprietary schemas Limited external supplier visibility ERP and MES identifiers differ in most cases Difficult to adapt for event-level traceability
Best suited for	Plant-level and lot/batch traceability, where most relevant information already resides in ERP/MES
Example	A manufacturer traces Finished Lot F-105 back to Production Order 8821, Material Lots M-21/M-22, the production line, and recorded quality results

Table 2. The components of centralized ERP/MES architecture.

Centralized ERP/MES architecture remains one of the most practical choices for manufacturers whose primary requirement is plant-level or lot-level genealogy. Its main limitation appears in cases beyond transactional manufacturing records, where the need for high-frequency machine data rises.

IoT- and Historian-Centric Architecture

Thanks to IoT- and historian-centric architecture, PLCs, sensors, machines, industrial gateways, SCADA systems, OPC UA servers, and IoT platforms continuously collect information from the shop floor, and this data can be used at much higher temporal resolution than a traditional ERP.

For example, instead of knowing only that Product X was processed on Machine 4, the manufacturer may have second-by-second temperature, pressure, vibration, speed, energy, or torque data for the relevant production window, which are irreplaceable to power further AI capabilities.

Aspect	IoT- and Historian-Centric Architecture
Strengths	High temporal resolution Strong machine context with automated data capture Useful for predictive quality and condition monitoring Supports detailed process investigation High AI potential for predictive quality, anomaly detection, predictive maintenance, and process optimization
Weaknesses	Sensor observations alone don't establish business genealogy Gathers immense data volumes with potentially high storage costs Difficult long-term retention decisions Machine identifiers may not correspond with ERP/MES identities
Best suited for	Process monitoring, predictive quality, condition monitoring, equipment analysis, and high-resolution production history

Table 3. The components of IoT- and historian-centric architecture.

[The source draft listed “High AI potential…” under Weaknesses — moved here under Strengths, where the content belongs (repeating the same fix applied to the previous draft of this paper).]

Distributed-Ledger Architecture (DLA)

DLA offers a shifted view from a single manufacturer's database toward a shared or multiple-participant replicated record with selected events written to a ledger. This way, cryptographic mechanisms can detect further alteration and serve as a source of truth for all participants. The biggest drawback is that DLA doesn't check whether data was correct when originally captured. So in case of mistakes, cryptographic immutability preserves that incorrect information as well.

Aspect	Distributed-Ledger Architecture
Strengths	Trust and evidence shared records across organizations Potentially strong provenance
Weaknesses	Governance complexity Privacy and commercial-confidentiality concerns Integration cost Difficult correction procedures Data accuracy is under question
Best suited for	Multi-party supply chains where provenance, shared evidence, or tamper detection is particularly important

Table 4. The aspects of Distributed Ledger Architecture (DLA).

Hybrid Event Architecture

A hybrid event architecture combines existing operational systems with an integration and event layer. For instance, ERP remains authoritative for suppliers, orders, and inventory, while MES handles production execution and QMS retains quality records. Additionally, APIs with event streaming and master-data services connect the relevant databases into a single shared environment.

For example, a supplier lot gets an ERP receipt, then it goes to an MES work order with machine event tracking. Further, it's marked as a serialized product and undergoes QMS inspection. All that remains is to track warehouse shipment and delivery to a customer.

It's essential to note that this paper adopts the hybrid model as a front-runner, as it allows manufacturers to incrementally modernize traceability without replacing current functional platforms.

Aspect	Hybrid Event Architecture
Strengths	Combines business and machine context Preserves existing systems Supports event-level genealogy Scalable across plants and suppliers Suitable for multiple compliance regimes and standards such as OPC UA, ISA-95, and EPCIS Strong foundation for analytics and AI
Weaknesses	Greater initial integration effort Requires strong identity and master-data management Need for event schemas governance Complex across legacy environments Requires clear ownership across IT, OT, quality, and data teams
Best suited for	Enterprise and multi-site traceability, regulatory reporting, cross-enterprise genealogy

Table 5. The aspects of hybrid event architecture.

To sum up, there is no universally “best” architecture. The choice is more about selecting an architecture according to the required scope. ERP/MES-centric systems remain effective for transactional genealogy; IoT-centric systems excel at capturing production conditions; DLA can strengthen multi-party provenance; and hybrid event architectures are the most flexible when it comes to operational traceability, including AI.

2.4 Manufacturing Traceability with AI

AI largely invests in traceability if records contain enough context to transform patterns into signals. 4 main groups are particularly valuable for manufacturers:

Anomaly detection. AI identifies unusual combinations and flags these events, answering the question “Which products have genealogy patterns statistically similar to previously defective products?”

Predictive quality. Quality models connect process parameters and genealogy with downstream inspection results through the chain from material, supplier, and machine, via tools and process parameters, to inspection and final quality outcome; however, without that contextual chain, the accuracy may be misleading.

Computer vision paired with genealogy. Computer vision linked to the corresponding unit or batch becomes one event in the product history.

Generative AI for traceability investigation. LLMs and retrieval-based systems offer natural-language interfaces for users, accelerating the investigation. In particular, generative AI assists in summaries and draft documentation; however, the final approval should be left to human controls.

2.5 Research Gap

Although existing manufacturing literature and standards address enterprise-control integration, serialization, supply-chain events, and so on, they frequently treat these concerns independently. However, a unified architectural question often remains unattended. How can the same traceability support operational results (genealogy, compliance, cross-enterprise interoperability, and AI analytics)?

A second gap is about maturity assessment. Digitization doesn't equal AI readiness; this way, mere digital transformation lacks consistent identifiers or governance. Therefore, the following sections contain the proposed reference architecture and maturity model that eliminate these gaps and offer compliance and AI readiness as one high-quality traceability infrastructure.

3. Recommended Traceability Architecture

3.1 Design Requirements

Before any specific recommendations, the preceding analysis shows the necessary architecture requirements for the best effect of the initiative:

- **Granular identity.** The architecture should support appropriate combinations of identities for materials, batch, serialized items, production orders, equipment, suppliers, facility, and shipment/containers.
- **Event-level genealogy.** The architecture must reconstruct transformations rather than merely store the current state. A useful event should answer the following questions:
  - What object was involved?
  - When did the event occur?
  - Where did it occur?
  - Why did it occur?
  - How was it processed?
  - Who/what executed or approved it?
  - What was the outcome?
- **Auditability.** Provenance, modification history, timestamps, responsible actors, and evidence relationships must be written in the architecture to ensure compliance.
- **Interoperability.** In most cases, production requires multiple locations and vendors. That's why traceability should include vendors' proprietary identifiers or schemas.
- **AI data readiness.** As a baseline, the architecture must save context for analytical use, such as labels and metadata for historical and predictive observations.
- **Jurisdiction-aware reporting.** Compliance outputs should therefore be generated from common traceability information, not remain as independent data silos.

3.2 Reference Architecture Overview

As a result of the aforesaid, the proposed architecture contains several logical layers depending on the traceability scope:

Layer 1. Data Capture and Edge. This layer reflects the physical production reality, adding edge gateways without modern interfaces. It can normalize older industrial protocols without immediate equipment replacement.

> *[Layer 2 (Integration and Normalization) is named in the abstract and contributions but its descriptive paragraph is still missing from the source draft, as in the prior version of this paper — add it here before publishing. Table 6 below (labeled as Layer 3 content) draws on components — APIs, message brokers, event streaming, master-data services — that are naturally part of this Layer 2 description.]*

Layer 3. Traceability System of Record. This layer inserts durable genealogy. The table below describes the components and work principles involved.

Component	Role	Example
OPC UA Gateways	Connect industrial equipment and expose machine data through standardized, structured	Read temperature, pressure, machine state, and cycle data from PLCs and

Component	Role	Example
	interfaces	associate them with Machine M17
APIs	Exchange traceability information between ERP, MES, QMS, WMS, supplier platforms, and external systems	MES retrieves Work Order WO8821 and material information from ERP
Message Brokers	Reliably route messages and events between systems without requiring direct point-to-point connections	A completed production cycle generates an event consumed by MES, QMS, and the traceability platform
Event Streaming	Processes continuous streams of production and supply chain events, often in near real time	Stream machine events, scans, inspections, and material movements as production occurs
ETL/ELT Pipelines	Extract, transform, normalize, and load historical or batch data from multiple systems	Combine historical ERP orders, MES production records, and QMS inspection results for genealogy analysis
Master-Data Services	Maintain authoritative definitions for products, materials, suppliers, equipment, locations, and other entities	Ensure “M17,” “Machine-17,” and ERP asset “00017” resolve to the same production asset
Identity Resolution	Connect identifiers referring to the same physical or business entity across different systems	Link Supplier Lot A91, internal material lot ML-4582, and the corresponding received shipment
Schema Validation	Verify that incoming events contain required fields, formats, units, and identifiers before they enter the traceability record	Reject or quarantine a temperature event that lacks a machine ID, timestamp, or valid unit of measurement
ISA-95 Mappings	Map data to a consistent manufacturing hierarchy and terminology across enterprise and shop-floor systems	Associate Machine M17 with Heat Treatment Cell 2 to Line 4 to Production Area A to Plant 01
EPCIS Event Transformation	Convert operational data into standardized traceability/visibility events suitable for cross-system or cross-enterprise exchange	Transform a warehouse scan into an EPCIS event describing what product moved, when, where, and why

Table 6. The components of the Traceability System of Record, Layer 3, in the proposed architecture overview.

Layer 4. Analytics and AI. Mainly, AI applications should reference the identities established in preceding Layers 2 and 3. An AI prediction then becomes traceable based on the correlation between model versioning/input data and prediction and confidence. This way, you add traceability of AI decisions, not vice versa, merely AI bolt-on traceability. The following table gives examples of how components work in the system.

Component	Role	Example
Feature Pipelines	Transform production and traceability records into consistent variables	Convert machine temperature into a feature for a quality-risk model
Analytical Warehouses/Lakehouses	Consolidate data for analysis	Combine two years of MES, ERP, QMS, historian, and supplier data to investigate recurring defects
Model Training	Uses historical outcomes to develop predictive models	Train a model to predict the probability of a defective component

Component	Role	Example
Model Registry	Maintains controlled versions of models	Record that Quality Model v3.4 was approved on a specific date and replaced v3.3 for Production Line 4
Predictive Quality Models	Estimate the likelihood of defects	Flag Serial 009173 as high-risk because its material resembles previous defective units
Anomaly Detection	Identifies unusual production, material, supplier, or supply-chain patterns	Detects an unexpected combination of supplier lot and unusually high rework frequency
Computer Vision	Analyzes images or video for automated quality inspection	Detects a surface defect
RAG/LLM Use Cases	Allows users to query governed manufacturing through natural language	A quality engineer asks, "Which products may have been affected by the Line 3 temperature excursion?" and receives an evidence-linked response
Model Monitoring	Tracks model accuracy and drift after deployment	Detects that predictive-quality accuracy has declined

Table 7. The components of AI analytics in Layer 4 of the proposed traceability architecture.

Layer 5. Compliance and Reporting. This fifth, top layer translates common traceability information into specific evidence products. The next table shows the possible components:

- **DPP Records** include product identity, materials, sustainability attributes, manufacturer information, lifecycle information, and relevant compliance data.
- **CBAM Evidence Packages** include goods identity, production facility, production quantity, direct/indirect emissions data where applicable, calculation inputs, and source evidence.
- **Recall Reports** include serial numbers, material genealogy, production dates, equipment, inspection results, shipments, and customer/distribution destinations.
- **FDA Traceability Exports** include Traceability Lot Codes, Critical Tracking Events (CTEs), and Key Data Elements (KDEs) with transformation and shipping information.
- **UDI-Related Records** include device identifier, production identifiers, lot/batch, and serial number.
- **Customer Genealogy Reports** include material lots, suppliers where appropriate, production order, manufacturing steps, inspection results, serial numbers, and shipment information.
- **Certificates of Analysis (CoA)** include batch/lot identity, test methods, measured values, specification limits, laboratory results, approval.
- **Sustainability Disclosures** include energy use, emissions, material composition, recycled content, and lifecycle information.
- **Supplier Provenance Reports** include supplier identity, location, material origin, certificates, shipment, and receipt events.
- **Auditor Evidence Packages** include source records, timestamps, approvals, audit logs, calculation methodology, document versions, and data lineage.

All layers unite into a unified, evidence-driven architecture. Machine data receives its context and connects to product genealogy, enabling further AI-supported analysis for relevant compliance audits. Importantly, this architecture doesn't require manufacturers to replace existing legacy systems and equipment; instead, it creates a common traceability layer across them.

Figure 1. Proposed Logical Architecture.

3.3 Data Standards and Interoperability

Neither standard should be interpreted as the complete architecture, nor can we propose a complete toolkit. Interoperability emerges from their combination. In particular, OPC UA describes a machine but doesn't replace genealogy; ISA-95 structures manufacturing context but doesn't replace product identifiers; EPCIS represents visibility events but doesn't replace MES execution; DPP defines external information requirements but doesn't itself capture production events.

For further illustration, the following table represents the possible standards and their roles in the architecture.

Standard	Purpose	Architectural role
ISA-95/IEC 62264	Enterprise/manufacturing integration	Defines operational context and boundaries between enterprise and

Standard	Purpose	Architectural role
		manufacturing functions
OPC UA	Industrial interoperability	Connects equipment, control systems, gateways, and software using standardized information models
GS1 identifiers	Globally standardized identification	Product, location, logistics, and trading-party identity
GS1 EPCIS 2.0	Visibility/event information	Exchange and query of supply chain and traceability events
GS1 Digital Link	Identifier-to-digital-resource connection	Connects product identifiers with online information/services
EU DPP framework	Product sustainability and lifecycle information	External compliance/product-information layer
ISO 9001 / sector QMS frameworks	Quality-system controls	Identification, controlled records, quality evidence
IATF 16949	Automotive quality management	Automotive production and supplier traceability context

Table 8. Standards and their architectural roles.

3.4 Security, Integrity, and Access Control

As traceability creates a gap between security and confidentiality, so supply-chain participants need to prove provenance without exposing sensitive production information. The architecture itself must contain safeguards for this case, for instance, differentiated access as follows:

- **Role- and attribute-based authorization.** Access may depend on role, product, facility, contractual relationship, etc.
- **Tamper evidence.** Important events can be protected through event signing, cryptographic hashes, trusted timestamps, etc.
- **Data lineage.** Every reported value should be traceable back to its origin for an auditor to evaluate not only the final number but also the evidence chain behind it.
- **Segregation of operational and external data.** External information should be published through controlled APIs or reporting services with explicit schemas and access rules.

4. Compliance Mapping

4.1 Regulatory Drivers Summary

The following table proves that the underlying evidence overlaps. This way, manufacturers should avoid one independent traceability database for each regulation but have many at the same time.

Framework	Status	Traceability implication	Layers
EU ESPR / Digital Product Passport	Framework in force; detailed requirements depend on product-specific measures	Persistent identity, structured/interoperable product information, controlled access	2, 3, 5
EU CBAM	The definitive regime applies	Embedded-emissions evidence	1–3, 5

Framework	Status	Traceability implication	Layers
	from Jan. 1, 2026	associated with covered imports and production	
FDA UDI	Implemented framework	Persistent medical-device identification and production attributes	2, 3, 5
EU MDR UDI	In force	Device identification and traceability	2, 3, 5
IATF 16949 with OEM requirements	Current automotive quality framework	Product genealogy and controlled production	1–3
EU CSDDD	Framework amended; implementation timeline evolving	Value-chain due diligence and supporting evidence	2, 3, 5

Table 9. Selected regulatory drivers and architectural implications as of August 2026.

4.2 Architecture-to-Requirement Mapping

The main principle is that one traceability record should satisfy multiple downstream requirements. For instance, product identity established in Layer 2 can't be recreated separately for recalls and so on. Instead, the architecture creates the identity and evidence chain once and exposes the appropriate parts to each compliance or operational use case. The following table connects business requirements and specific architecture components.

Traceability Capability	Requirement	Architecture Response	Layers	Example
Product Identification	Establish persistent identity	Identity resolution and master-data services	Layer 2-3	Serial 009173 is linked to its product model and genealogy
Production Genealogy	Reconstruct which materials, processes, facilities, equipment, and operations were involved	Event capture combined with an event store and genealogy graph	Layer 1-3	A defective assembly is traced back to the supplier, machine, and inspection with designated treatment
Provenance	Establish where materials/products originated and transferred	Supplier master data, shipment, and supporting documents	Layer 2-3	Raw material is traced from Supplier A via shipment to Plant B for Distributor C
Environmental Evidence	Associate emissions and energy consumption with relevant production	Sensor capture, facility and process mapping, calculation rules	Layer 1-5	Electricity consumption from a production line is associated with the relevant production quantity

Table 10. The connection between business and compliance requirements in the traceability architecture.

5. AI Readiness Assessment

Instead of training models on isolated MES, ERP, or sensor datasets, manufacturers can use complete production histories connected to their operational context. However, having large volumes of manufacturing data doesn't automatically turn a factory into an AI-ready one.

AI readiness depends on complete, consistently identified, historically available, granular, and traceable data. The following sections therefore examine the use cases enabled by a mature traceability architecture and the operational capabilities required to deploy them.

5.1 AI Use Cases of Traceability in Manufacturing

The most reliable applications of AI in manufacturing traceability include the following cases:

- **Predictive quality scoring.** The system combines genealogy and process parameters to estimate the probability of non-conformance before final inspection.
- **Supply-chain anomaly detection.** Models identify unusual patterns and flag them for further action.
- **Root-cause analysis.** Graph analytics and machine learning automatically rank correlations between defects and upstream factors for further analysis.
- **AI-assisted compliance investigation.** A governed LLM interface retrieves relevant records and generates its answer with references to the records.
- **Automated draft reporting.** Generative AI transforms structured evidence into first-draft narratives for all kinds of reports.

Furthermore, in all cases, human review remains obligatory to some extent, especially where decisions bring potential risks such as regulatory compliance or safety.

5.2 Data Readiness Criteria

Returning to the differences between accuracy and traceability, data requires more than volume. The following requirements should be applied to data to consider it reliable:

- **Completeness** — gaps can invalidate otherwise sophisticated analysis.
- **Identity consistency** — whether the same data can be reliably recognized across systems.
- **Temporal quality** — proving the precision to reconstruct the sequence.
- **Context** — to explain what a sensor measurement or transaction represents.
- **Outcome labeling** — for data to be linked to outcomes.
- **Lineage** — to trace data to its original record.
- **Representativeness** — to confirm the data represents sufficient historical information.
- **Governance** — to check who controls data.

All considered, we can describe the separate maturity levels for a manufacturing environment before any further action can be taken.

5.3 AI Readiness Maturity Model

This paper proposes the structure of a manufacturing traceability AI readiness maturity model with 5 levels explained in the table below.

Level	State	Data	Architecture	AI capability	Governance
0. Manual	Paper/spreadsheets	Fragmented, manually entered	No common traceability architecture	Research only	Individual ownership
1. Digitized	Key records are electronic	ERP/MES silos	Point integrations	Local analytics	System-specific controls
2. Integrated	Cross-system genealogy	Normalized identities and timestamps	APIs/events	Reliable BI and initial ML	Enterprise data ownership

Level	State	Data	Architecture	AI capability	Governance
3. AI-Augmented	Contextual event history	Labeled, lineage-aware datasets	Data and feature pipelines under MLOps	Predictive quality, anomaly detection, AI investigation	Model governance
4. Adaptive with Oversight	Near-real-time closed information loop	Continuously validated data	Event-driven traceability and governed AI	Automated recommendations	Continuous monitoring

Table 11. Manufacturing Traceability AI Readiness Maturity Model.

Note that the final goal is not autonomous AI but controlled autonomy with traceable oversight.

5.4 Risks and Governance

Moving AI into production introduces an additional challenge to eliminate risks. This way, the governance should be designed as part of the AI architecture rather than added after deployment. Each major risk should have corresponding control measures that define how to detect and mitigate it.

Risk	Why It Matters	How to Address It
False Correlations	Models may associate defects with a supplier, machine, or process incorrectly	Validate predictions against engineering knowledge and investigate confounding variables with human confirmation before consequential actions
Model Drift	As operating conditions change, historical models become less reliable	Continuously monitor model performance and data distributions with installed drift thresholds. Ground outputs in controlled sources via RAG for model hallucination monitoring
Automation Bias	Operators may accept AI recommendations without sufficient independent verification	Display confidence and supporting evidence and require human review for critical decisions
Explainability	Regulators, quality teams, or engineers may need to understand why a recommendation was produced	Preserve input features, model version, and relevant evidence
Data Leakage	Traceability data may expose supplier-sensitive data	Apply encryption, logging, and least-privilege access

Table 12. AI risk categories and governance controls.

6. Conclusion

The principal architectural challenge of manufacturing traceability is preserving the identities, relationships, events, semantics, and lineage according to what happened to a product across its lifecycle in reality. Consequently, this paper proposed a 5-layer architecture connecting physical data capture to integration, durable genealogy, AI, and compliance reporting.

The central conclusion proves that both compliance readiness and AI readiness increasingly depend on trustworthy contextualized data. Future research should validate the framework through production deployments and develop quantitative maturity benchmarks with a recommendation framework, as that question may define the next generation of industrial traceability.

References

European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1781 establishing a framework for the setting of ecodesign requirements for sustainable products. Official Journal of the European Union.

European Parliament and Council of the European Union. (2023). Regulation (EU) 2023/956 establishing a carbon border adjustment mechanism. Official Journal of the European Union.

European Parliament and Council of the European Union. (2017). Regulation (EU) 2017/745 on medical devices. Official Journal of the European Union.

European Parliament and Council of the European Union. (2024). Directive (EU) 2024/1760 on corporate sustainability due diligence. Official Journal of the European Union.

European Parliament and Council of the European Union. (2025). Directive (EU) 2025/794 amending the dates of application of certain sustainability reporting and due diligence requirements. Official Journal of the European Union.

European Commission. (2026). Carbon Border Adjustment Mechanism: Definitive regime.

GS1. (2022). EPCIS and Core Business Vocabulary (CBV), Version 2.0.

International Society of Automation. ANSI/ISA-95 / IEC 62264: Enterprise-control system integration.

OPC Foundation. OPC Unified Architecture specifications and companion specifications.

International Organization for Standardization. ISO 9001: Quality management systems — Requirements.

International Automotive Task Force. IATF 16949: Quality management system requirements for automotive production and relevant service parts organizations.

Appendix A. Glossary of Terms and Acronyms

- AI — Artificial Intelligence
- CBAM — Carbon Border Adjustment Mechanism
- CTE — Critical Tracking Event
- DLT — Distributed Ledger Technology
- DLA — Distributed Ledger Architecture
- DPP — Digital Product Passport
- EPCIS — Electronic Product Code Information Services
- ERP — Enterprise Resource Planning
- ESPR — Ecodesign for Sustainable Products Regulation
- HITL — Human in the Loop
- IoT — Internet of Things
- KDE — Key Data Element
- LLM — Large Language Model
- MDM — Master Data Management
- MES — Manufacturing Execution System
- MLOps — Machine Learning Operations
- OPC UA — OPC Unified Architecture
- PLC — Programmable Logic Controller
- QMS — Quality Management System

- RAG — Retrieval-Augmented Generation
- RFID — Radio-Frequency Identification
- UDI — Unique Device Identification
- WMS — Warehouse Management System

Appendix B. Standards and Regulatory Quick Reference

Framework	Issuing body	Domain	Primary relevance to traceability	Status considered in this paper
ESPR/DPP	EU	Products/ sustainability	Digital identity and structured lifecycle/product information	Framework in force
CBAM	EU	Carbon/imports	Embedded-emissions evidence	Definitive regime operating in 2026
CSDDD	EU	Supply-chain due diligence	Value-chain risk/evidence	In force; timetable amended
EU MDR UDI	EU	Medical devices	Device identification and traceability	In force
UDI	US	Medical devices	Device/production identification	Implemented
ISA-95 / IEC 62264	ISA/IEC	Manufacturing	Enterprise-control integration	Established standard family
OPC UA	OPC Foundation / IEC	Industrial systems	Secure semantic machine/system interoperability	Established and evolving
ISO 9001	ISO	Quality	Quality-management and identification controls	Established QMS standard
IATF 16949	IATF	Automotive	Automotive QMS and traceability	Current automotive QMS framework