International Journal of Computer Trends and Technology
ISSN: 2231–2803 / https://doi.org/10.14445/22312803/IJCTT-V74I8P103
Volume 74 Issue 8, 18-27, August 2026

Original Article

Supply Chain Visibility and End-to-End Traceability: Enhancing Food Safety and Consumer Trust in the Consumer Products Industry

Asha Kiran Ganesna

Chief Architect, MN, USA.

Received: 23 June 2026 Revised: 28 July 2026 Accepted: 16 August 2026 Published: 30 August 2026

Abstract *- The consumer products industry has prioritized food safety incidents and opaque supply chains, as they are no longer an afterthought for them; it's become a strategic priority. This article discusses the practical benefits for foodborne illness risk reduction, quicker recalls, and consumer-trust restoration that can be achieved through end-to-end supply chain visibility, built on the Internet of Things, distributed ledger technology, and the power of artificial intelligence and machine learning. It suggests a 5-layer reference architecture covering data capture, connectivity, data integration and trust, AI/ML analytics and end-user visibility applications, and outlines AI/ML use cases such as anomaly detection, predictive quality scoring, computer-vision quality control, and graph-based trace-back, plus a "workflow" spanning incident to recall enabling continuous model retraining. The article also examines obstacles to implementation, such as data standardization, integration with legacy systems, and cost, and ends with a vision of digital twins, generative AI, and predictive recall. The tables and diagrams illustrate a comprehensive analysis to provide a structured and evidence-based reference for practitioners and researchers on designing and evaluating modern traceability programs.*

Keywords *- AI/ML Analytics, Consumer Trust, EPCIS/Blockchain, Food Safety, Supply Chain Traceability.*

1. Introduction

Consumer products supply chains are complex, multi-stakeholder, and multi-faceted networks that include various layers and are global in extent, especially for food and perishable goods. The ingredients in a packaged product can be from a dozen countries, go through multiple co-packers and distribution centers, and then arrive at thousands of retail stores before finally ending up on a consumer's plate. This complexity allows for cost savings and product diversity, but also creates numerous opportunities for product contamination, fraud, temperature abuse, and documentation errors. In the event of a failure, the question that seems so simple to ask is actually quite hard to answer—the question is, which lots, which locations, and how fast can they be isolated?

Supply chain visibility is the ability for an organization and its partners to see when, where, and how goods are traveling in the supply chain exactly as they were intended [1]. End-to-end traceability takes this logic one step further, creating a timeline of every product from raw material origin to its manufacturing, distribution, and retailing processes and even to the individual consumer, where possible [2]. Traceability is different from visibility, as it can be asked, "Where has it been and what happened to it along the way?" Combined, they constitute the backbone of today's food safety management and of confidence in packaged food and consumer product brands that grace the consumer.

Although there have been multiple studies about food-supply-chain traceability, they all focus on one aspect of the system, and several existing studies examine either IoT sensing, blockchain/EPCIS data sharing, or AI/ML analytics separately and do not propose a system of these aspects. Evaluations of IoT-based cold-chain monitoring and applications focus almost exclusively on the monitoring aspect, with little focus on how the resulting event data can inform a shared trust layer by all parties; studies on Blockchain traceability tend to concentrate on the chain's ability to provide information on data provenance and offer minimal insight into what happens after data provenance arrives, and forms part of a controlled, auditable, real-time recall. Evaluations of IoT-based cold-chain monitoring tend to focus almost entirely on the monitoring part of the system, and studies of AI/ML applications for anomaly detection or computer-vision inspection are generally evaluated in isolation, with the question of how the output of this monitoring feeds into a controlled, auditable real-time recall sequence being insufficiently explored. If there is no single source for practitioners to follow to coordinate these layers end-to-end from sensor to shelf to consumer notification, or a common vocabulary for describing and comparing programs across organizations, then how can change be measured? It aims to close those gaps directly, proposing a five-layer reference architecture that combines data capture, connectivity, trust, analytics and visibility into the aforementioned continuum; outlining a variety of AI/ML use cases onto the proposed reference architecture; and providing a single incident end-to-end complete with a detection-to-recall flow, which can provide a unified, structured basis for designing and evaluating traceability programs which is not a set of disconnected technology reviews.

This article reviews the state of the art and the technology being developed to suggest a systematic approach to thinking about traceability architecture. It emphasizes the importance of artificial intelligence and machine learning, as they transform traceability systems from mere record-keeping tools to proactive risk management systems. The discussion centers on a reference architecture, the set of AI/ML use cases, an incident detection and recall workflow, and an honest account of the organizational and technical hurdles between the end-to-end principle of traceability and its realization.

The analysis is based on publicly available regulatory guidance, standards documentation, and industry practice and is designed to be a structured reference for researchers, quality and supply chain practitioners, and technology leaders considering, researching, or designing a traceability program. When the capabilities discussed exist in mainstream production use, for example, in the area of event sharing and cold-chain anomaly detection, the article distinguishes these capabilities from capabilities that are emerging or early-stage, such as graph-based risk propagation modeling and generative AI-assisted trace-back querying, so that the reader can set their scope of work accordingly.

2. The Food Safety and Consumer Trust Imperative

Until now, public health justification for traceability is robust. Foodborne illness is a constant threat in even the most advanced food systems, and leafy greens, spices, infant formula, and processed proteins have all been tied to outbreaks, showing how slow a disjointed, paper-based recordkeeping system can be to respond [3]. Performing a recall can take days or weeks for investigators to attempt to reconstruct supply routes via phone and fax before the recall can be scoped. By this time, the contaminated product is in the store and in consumers' hands. Every day of delay in a trace-back investigation is a day for continued, unnecessary exposure.

In turn, regulators have issued robust prescriptive recordkeeping requirements. Under Section 204 of the Food Safety Modernization Act (FSMA), the Food and Drug Administration (FDA) has issued the Food Traceability Rule for the United States, which requires standardized Key Data Elements to be maintained and communicated at each Critical Tracking Event (CTE) by firms that manufacture, process, pack, or hold foods subject to the FSMA on the Food Traceability List and a sortable electronic record, upon request, to be made available within 24 hours [4], [5]. In consultation with industry, the rule was originally proposed to go into effect on January 1, 2026; however, it has now been extended to become effective on July 20, 2028, due to industry concerns regarding the coordination requirements of a rule that only provides full public-health benefits after everyone in a supply chain is compliant. The extension does not diminish the underlying requirements; rather, it accepts that the requirement of end-to-end traceability is a multi-party data interoperability requirement and is not simply an IT project with a single firm. Similar systems are in place in other parts of the world, like the EU food information and traceability regulations, and an increasing number of national and regional regulations throughout the Asia-Pacific and Latin America.

Traceability is now a commercial differentiator, in addition to regulation [6]–[8]. Brands are trusted more by consumers than claims or marketing copy—and this is especially the case for younger generations, who are demanding evidence of origin, sustainability, and ethical sourcing. Retailers have followed suit by advancing their expectations through supplier scorecards, electronic advance shipment notices, and GS1-compliant labeling requirements further down their supply chain [9]. The overall impact is that there are now three equally compelling forces at play: regulatory compliance, operational risk management, and brand trust—all of which are driving an increasing need to think of traceability as a strategic asset, rather than a cost center.

3. Defining Supply Chain Visibility and End-to-End Traceability

Often, traceability is done by default as "one-up, one-down" record-keeping, where each party in the supply chain only keeps a record of the party up and down in the chain [12]. This option will meet the minimum record-keeping requirements set by the regulatory authorities, but no one, not even the brand owner, will have a total picture of the chain. To fully trace back, therefore, one would have to contact all parties of interest in turn, which is slow, requires manual cooperation, and would be susceptible to loss of records or inaccuracies at any step.

End-to-end traceability, on the other hand, seeks to establish a permanent, shared document following the whole chain of custody, usually based on the GS1 Electronic Product Code Information Services (EPCIS) standard. EPCIS establishes a common language of “what, where, when, and why" activities that can be created and shared between systems and organizations consistently [10], [11]. When combined with a standardized party, product, and location identifier, an EPCIS event enables computational resolution of a traceback query, eliminating what was once a days-long manual process into a minutes-long computer-based operation. There is a relationship between the two ability elements, visibility and traceability, but there is a difference. The main challenge with visibility is about real-time or near-real-time situational awareness: Where is this shipment? What is its temperature? Is it on time? Traceability is a historical and forensic concept: If you already have a finished product or a raw material lot, then traceability must be able to describe all the places it has been and what it has touched. A mature architecture can provide both, leveraging the same event data for offering real-time dashboards and post-event investigations.

4. A Reference Architecture for End-to-End Traceability

To make the idea of end-to-end traceability a reality, a system has to be able to collect the data at the point it originates, transport it reliably through organizational and network boundaries, store it in a format that can be relied on by multiple parties, analyze it to look for risk/quality signals, and report it out to the people and systems who must take action based on it. The five-layer reference architecture shown in Figure 1 is representative of existing industry best practices of leading consumer products and food manufacturing companies.

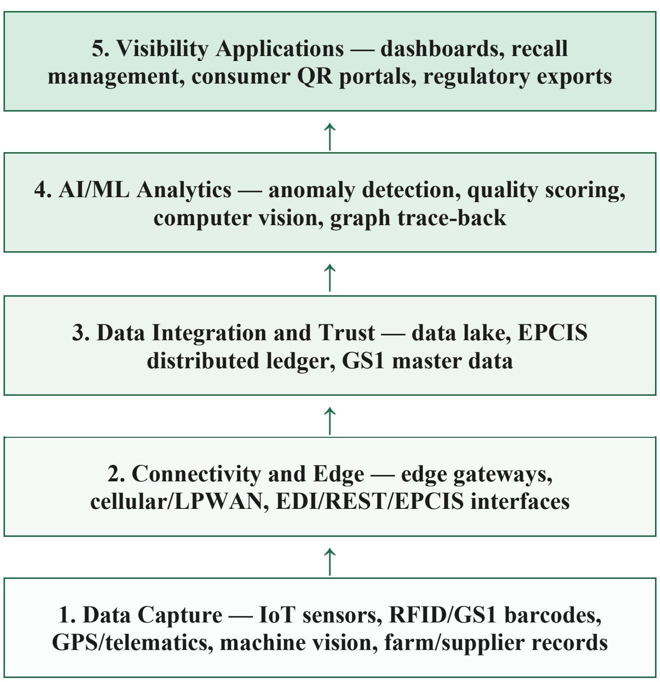

Fig. 1 Layered reference architecture integrating IoT capture, connectivity, ledger-based data trust, AI/ML analytics, and end-user visibility applications

The data capture layer is the bottom layer of the architecture. It includes IoT sensors that record temperature, humidity, and shock exposure; RFID tags and GS1 standards' barcodes with batch/lot identifiers; GPS and fleet telematics on cold chain-carrying vehicles; machine-vision cameras positioned along production lines; and enterprise systems at supplier and farm levels that document inputs, certifications, and growing/manufacturing conditions.

The connectivity and edge layer normalize and move this data through the use of edge gateways, which translate protocol information and buffer locally, and through cellular and low-power wide-area networks that provide connectivity to the field and are used for various applications, including the EDI, REST, and EPCIS dialects that different trading partners require.

The data integration and trust layer is the space where the individual data streams turn into a shared record that's trusted. The larger body of structured and unstructured supporting data can be stored in a cloud-based data lake or data warehouse; a distributed ledger and an immutable, timestamped log of EPCIS events can be used by multiple parties to query without having to trust a single central authority; and a master data management system, usually rooted in a GS1 registry of product, location, and party identifiers, can ensure the same meaning of the lot code or GLN to all systems consuming it.

It is important to integrate with existing ERP, warehouse management, and transportation management systems, as traceability data is useless if it is in a system separate from daily operations.

The AI/ML analytics layer, which will be discussed in detail in Section 5, converts this data into risk intelligence by recognizing abnormal temperature variations, calculating the shelf life and quality of products, identifying defects or contamination on production lines, and modeling known quality events & how they impact the broader supplier and distribution network.

Lastly, it displays and utilizes this intelligence in front of individuals and systems that need it, using real-time dashboards with track and trace data for supply chain and quality teams, automated recall management, consumer-facing QR code portals with origin and freshness data, and structured exports for regulatory and audit reporting.

4.1. Data Capture Technology Comparison

Due to the variety of traceability needs and the fact that no single technology can meet all the demands, most deployments involve multiple technologies from the list provided in Table 1, depending on the product, risk, and budget of the segment of the supply chain being instrumented.

Table 1. Comparison of core data capture technologies used in end-to-end traceability architectures

Technology	Primary use	Data captured	Typical Latency	Key Limitation
IoT temperature/humidity loggers	Cold-chain monitoring	Continuous environmental readings	Seconds to minutes (real-time)	Battery life; sensor placement bias
RFID / GS1 barcodes	Lot and unit identification	Batch, lot, and serial identifiers	Near-instant at scan point	Requires line-of-sight or read-zone discipline
GPS / fleet telematics	In-transit location tracking	Position, route, dwell time	Seconds to minutes	Signal loss in tunnels, dense urban areas
Machine vision cameras	Line-side quality inspection	Images/video of the product surface	Milliseconds (inference-dependent)	Sensitive to lighting and camera drift
EPCIS / blockchain event logs	Cross-party event sharing	Structured what/where/when/why events	Minutes (batch or streaming)	Requires partner-wide standard adoption

5. The Role of Artificial Intelligence and Machine Learning

Traditional traceability systems are primarily based on an "after the event is reported" approach, usually triggered by a regulator, retailer, or consumer complaint. AI and ML change the system's focus — from reactive to proactive — by leveraging the same set of events and sensors already collected to identify problems sooner, make the response effort more manageable, and improve the accuracy of identifying problems as more events are perceived over time. There are specific categories of AI/ML applications with a significant impact on food and consumer product traceability.

Anomaly detection models are usually trained using one of the isolation forests, one-class classifiers, and Recurrent Neural Network (RNN) autoencoders, and these models continually score incoming sensor streams for patterns of "normal behavior" [13]. The underlying models for anomaly detection are typically built using one of the following algorithms: one-class classifiers, Recurrent Neural Network (RNN) autoencoders, or Isolation Forest, and they continuously score incoming sensor streams to see if they follow patterns typical of “normal" behavior.

These models can learn the normal temperature range for a specific route, vehicle, or type of product and alert only if a temperature condition outside the range is detected, such as if a refrigeration unit were degrading slowly rather than abruptly.

Instead of relying on fixed threshold alarms, which trigger too many false alarms in naturally variable cold-chain environments, these models learn the normal range of variation for a given route, vehicle, or type of product and alert only when something abnormal is taking place, such as a refrigeration unit gradually failing over time rather than abruptly. These models are known as predictive quality and shelf-life models and are typically based on a gradient-boosted tree architecture or regression ensemble of models and include microbial testing history along with time-temperature exposure and packaging characteristics to predict remaining shelf life or spoilage risk for a given lot before a quality failure occurs.

Typically, CNNs are trained on image collections annotated with quality grades; the computer vision models automatically perform high-throughput quality grading and foreign object and/or defect detection at line speed, identifying contamination and packaging defects that may be missed using human visual inspection or may only be detected on an intermittent basis if the inspector is tired [14].

Graph models, such as graph neural networks modeling the supplier-shipment relationship network, enable trace-back and trace-forward queries as well as risk propagation analysis—not only identifying the lot that was just reported as contaminated but also all downstream products, distribution centers, and retailing outlets that receive material associated with a contamination event, and even simulating how a hypothetical contamination event at a specific supplier node would likely propagate through the wider network before it happens.

Table 2 summarizes these use cases, along with the modeling techniques, sample input data, and business outcomes commonly associated with each use case, which offers a helpful guide for any organization embarking on an AI/ML-enabled traceability roadmap.

Table 2. AI/ML use-case matrix mapping techniques, data inputs, and business outcomes across the traceability program

Use case	Representative ML Technique	Primary Input Data	Business Outcome
Cold-chain anomaly detection	Isolation forest / LSTM autoencoder	Time-series temperature and humidity logs	Earlier detection of equipment failure or temperature abuse
Predictive quality and shelf-life scoring	Gradient boosting / regression ensembles	Time-temperature history, microbial test results	Proactive rerouting, markdown, or hold decisions
Computer vision quality control	Convolutional neural networks	Line-side camera images and video	Line-speed defect and contaminant detection
Graph-based trace-back / trace-forward	Graph neural networks	EPCIS events, shipment and supplier relationships	Faster, more complete recall scoping
Demand and risk-adjusted forecasting	Ensemble time-series models	Sales history, seasonality, supplier risk scores	Reduced overstock exposure to at-risk lots
Supplier document and certificate review	NLP / document classification	Certificates of analysis, audit reports	Faster onboarding, automated compliance flags

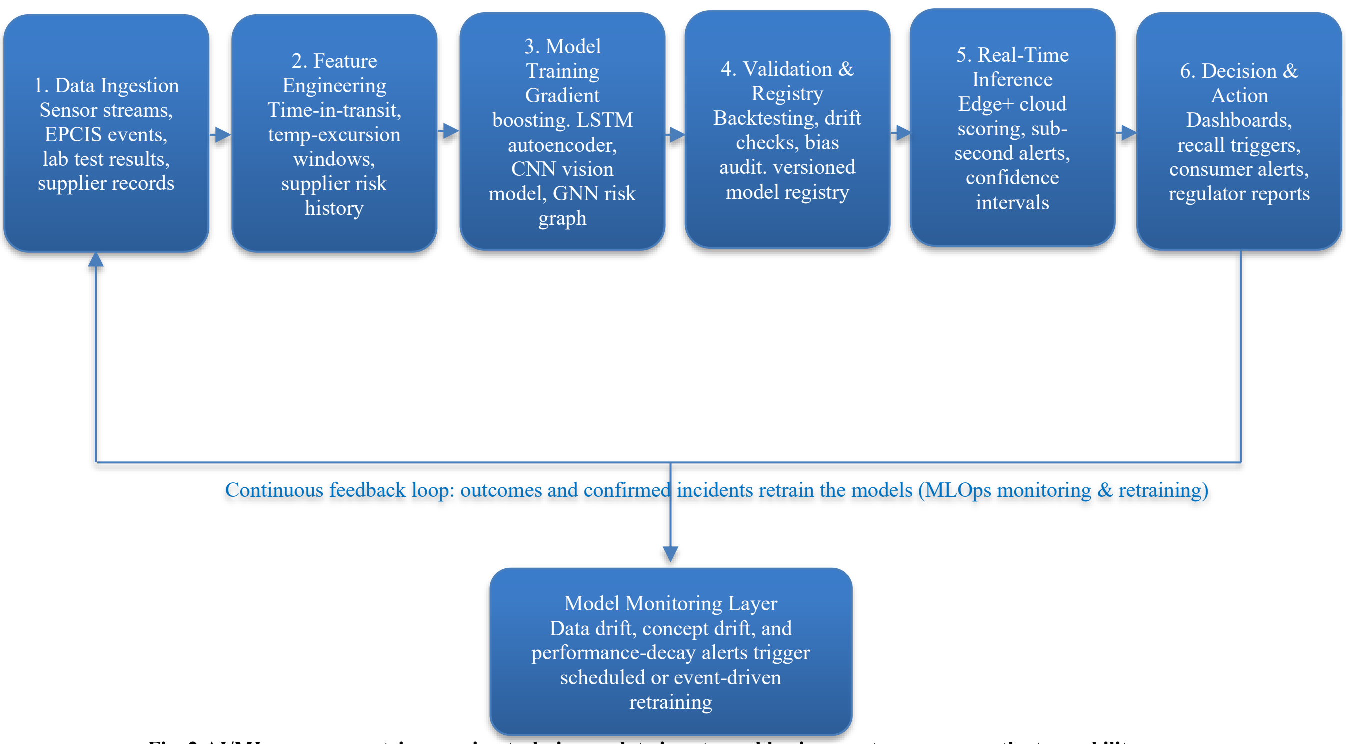

Fig. 2 AI/ML use-case matrix mapping techniques, data inputs, and business outcomes across the traceability program

Figure 2 shows the operational flow of these models—from ingestion of raw data, feature engineering, model training, model validation, real-time inference, decision layer, to action—the culmination of which results in dashboards, recall workflows, and consumer notifications.

The other equally relevant — although less obvious — piece of a production AI/ML traceability system is model governance [15]. Quality and regulatory affairs staff need to be able to interpret the results of the models, rather than just a score, for consequential actions like product holds, recall notification, etc.; for example, to see which specific anomaly (e.g., a sustained temperature excursion above/below a certain limit, for a certain duration) triggered a certain classification. Therefore, models must be designed to produce understandable results, not just a score. Documenting the model's validation criteria, including historical incidents against which it was tested, and keeping the model under test and the model's versions in a versioned model registry that tracks the operation of which model version resulted in a particular decision is becoming far more common and is expected to meet the needs of an internal audit and be the basis for any regulatory review after a model is recalled [16]. The model monitoring layer is perhaps one of the more important components and is underused in this pipeline; that is, the models are periodically retrained, or it can be event-driven when data drift and concept drift are detected or when there are changes to the statistical properties of the incoming data, or the underlying relationship between inputs and outcomes. Seasonal shifts in the sourcing of produce, as well as packaging and new supplier onboarding, all cause the model to shift, and a static model trained once will quietly deteriorate in its accuracy over time. A key trait of the more mature implementations is to treat the implementation of traceability AI/ML as an operational system, and not a one-off analytics initiative.

6. From Detection to Recall: An AI/ML-Enabled Incident Workflow

The use of the architecture and models above is best demonstrated when a true food safety incident arises in the end-to-end process. This workflow goes from continuous ingestion of sensors to the post-incident training of the model, as illustrated in Figure 3; what would typically have taken days to complete can be reduced to hours, or shorter, by automating the detection, classification, and trace-back functions.

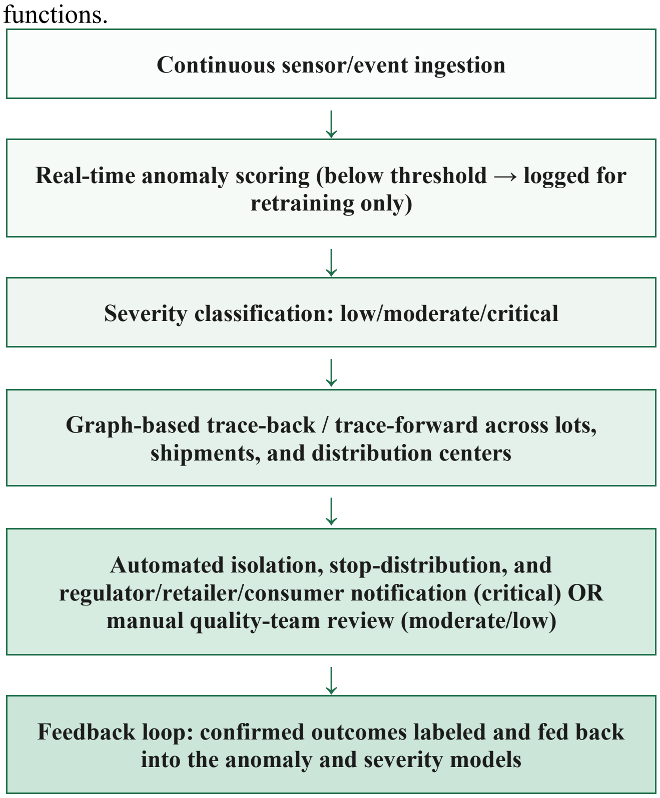

Fig. 3 AI/ML-driven workflow from continuous sensing to model retraining, closing the loop for continuous improvement

Events continuously ingested are events that are scored in real time by an anomaly detection model, which involves events that contain different types of information, such as temperature, location, and batch scan. Events that do not exceed the anomaly threshold are recorded to the data lake for continued monitoring and future training of the model without impacting the business flow, which keeps false positive rates below a threshold where quality teams start to ignore the system and develop “alert fatigue." If an event exceeds the threshold, then it is transferred to a severity classification model, which primarily uses the product category of the event (e.g., food, bulk material, or sugary confectionery), specific route or facility, and historical risk patterns of similar events to rank the likely event as low, moderate, or critical.

When an incident exceeds the moderate or higher level, a graph-based traceback and trace forward automatically gather all the linked traces of the lot, shipment, distribution center, or downstream retailer involved with the flagged incident, taking significantly less time to investigate manually and from a spreadsheet and reducing the likelihood of missing indirect connections. If severity determination exceeds a predetermined threshold or a predefined regulatory trigger, the system can automatically isolate affected lots throughout the network, stop additional distribution, and provide notifications to regulators, distribution centers, retail partners, and, if there is an application that allows it, affected consumers through multiple channels. Less severe incidents are transferred to a quality team for manual review and hold decisions—leaving human judgment for cases where automated action is not ready.

The concluding step in the workflow, and perhaps the most critical, is the feedback loop, where confirmed incidents and the resulting outcomes are labeled as training examples and added back into the models of anomalies and severity; this increases the accuracy of the detection process, and fewer false positives and false negatives are produced with each successive iteration. This closed-loop design is what makes a learning traceability system truly different from a static rules engine and why the use of AI/ML algorithms is an investment on an operational basis and not just for the sake of a one-time deployment.

7. Blockchain and Distributed Ledgers as a Trust Layer

An ever-returning barrier to the implementation of end-to-end traceability is not the lack of data, but the lack of trust between the multitude of independent organizations that will be required to provide the data. A retailer has little reason to take the self-reported quality record of a supplier on trust, and a supplier has every right to have concerns about sharing information that is commercially sensitive through a central database maintained by a single retailer. A problem being tackled by a group of technologies like permissioned blockchain is to solve it structurally, which is what multiple parties can write to and query a shared, cryptographically verifiable event log without any one party having control or the ability to change the log [17]. Most consumer product implementations take place on permissioned networks (not public networks) where access to participants in the supply chain is restricted to those who are vetted by the network, and visibility may be controlled by access rules, such as allowing a regulator to see full trace-back detail but limiting a competing retailer to the overall compliance status. Most often, the blockchain layer does not store bulk sensor data itself; rather, it stores the EPCIS-formatted summary of these events and the cryptographic hash of the underlying data stored in a traditional cloud data lake [18]. This hybrid approach combines the trust and tamper-proof features of a distributed ledger with the performance and cost-effectiveness of traditional cloud storage for the data itself.

It is interesting to note that blockchain is not a precondition for or a guarantee of good traceability. Similar results can be obtained by a well-governed, standards-based data sharing relationship on top of conventional databases in the context of a small, stable group of trading partners, and poorly integrated blockchain implementations can give the appearance of data immutability when the data actually isn't. The technology is most elegantly seen as a means to build multi-party trust at scale, particularly where several suppliers are linked in a many-to-many supply chain, and no party has the authority to serve as a trusted central authority.

8. Consumer Trust and Transparency

In-house traceability systems may be highly complex, but if they remain at the factory gate, they can only offer a small proportion of the benefit. Increasingly, consumer product brands are taking data out from their EPCIS streams and cascading it to consumer-facing applications, such as a QR code printed on food packaging that connects to a web portal that communicates details about the origin, farm or facility history, major certifications, and, in the case of perishables, freshness or harvest date information [19], [20].

The trust dividend that is gained from this transparency seems to happen on a few fronts. It helps substantiate origin and sustainability claims with verifiable facts instead of asking consumers to take marketing claims on faith; it provides some measure of confidence about the brand's operational competency; and it changes the nature of the recall experience for the consumer, because it enables a brand to identify quite specifically and immediately which particular product lots are involved, rather than issuing a general, imprecise recall message covering an entire product line, both avoiding unnecessary product disposal and helping prevent unnecessary harm to a brand's revenue and reputation.

Despite the ease of measurement, the real results of the transparency associated with consumer-facing brands remain methodologically challenging, as trust is a psychological construct that does not translate directly and uniformly into purchases, and brands should be careful not to ascribe increases in sales to transparency based on traceability alone if other marketing and pricing considerations are not taken into account. Nonetheless, engagement results, including how many QR codes were scanned, the length of time spent on the portal, and repeat scans for specific product batches, provide a useful, albeit limited, indicator of consumer interest; and some brands reported using the engagement patterns to decide how to prioritize the claims and data points to be shown most prominently in future packaging and portal redesigns.

This transparency comes with its dangers and difficulties. The structure of the supply chain can be exposed by product exposure and can show competitively sensitive supplier relationships; consumer-facing portals must be subject to the same care for product data quality—spotting an incorrect or stale origin claim by an involved consumer can have a more deleterious impact on consumer sentiment than any claim at all. Generally, the organizations that are interested in consumer-facing traceability can begin with a more limited number of verifiable, value-adding claims (like single sourcing or third-party certification) and then move to more broadened supply chain narratives.

9. Data Ethics and AI/ML Governance in Practice

While the technical benefits outlined above gain from extending traceability data into AI/ML models and Shared (Distributed) Ledgers, it creates ethical challenges, often in opposition, that go beyond the technical considerations. First, there is an uneven distribution of consent and data minimization throughout the chain – growers, small suppliers, and gig-economy drivers may not have firsthand knowledge of how their data is used, stored, and resold, and programs should only collect data for the specific purpose for which they are needed, not gather data because a sensor can. Second, an automated or model-assisted decision (such as a lot, a recall or a down-ranking of a supplier's risk score) has real commercial and reputation consequences, and therefore must be made by a named person, not a model version, and must be considered as a decision, not a model result. Third, supplier risk-scoring models trained on data from prior audits and incidents can risk being subject to risk coding bias against smaller or newer suppliers because they have less experience at being audited, and not because of substandard practices; therefore, periodically auditing risk scores for size- or tenure-correlated risk coding bias is warranted, not taken for granted, based on the assumption that a model developed on valid outcomes is automatically legitimate. Lastly, the immutability which makes a distributed ledger valuable for trust conflicts with the rights of data subjects to delete or amend sensitive data in the privacy regulations, such as the EU General Data Protection Regulation; in the current implementations, the solution is to store only the cryptographic hashes of the data, and the metadata of the events, without any personal or commercially sensitive information, on the distributed ledger itself, leaving the other data off-chain where it can be modified or deleted without impairing the chain of custody. All these things do not run counter to the adoption of AI/ML or blockchain; they are due to the fact that data governance is a design factor, as well as the capture of data and data analytics, and for this reason, it should be kept in mind throughout the design process, not as an add-on to deal with later, once a system is up and running.

10. Implementation Challenges

While it is obvious that regulatory and business needs and desires drive the end-to-end traceability program, it often fails to meet the program's original intent, and the reasons are more often organizational or data-oriented than technological. Table 3 provides an overview of the most frequently reported implementation challenges and strategies to mitigate them.

Table 3. Common implementation challenges in end-to-end traceability programs and corresponding mitigation strategies

Challenge	Description	Mitigation Strategy
Data standardization and interoperability	Trading partners use inconsistent identifiers, units, and event formats, preventing automated cross-party queries.	Adopt GS1 EPCIS and Core Business Vocabulary standards; require standard compliance in supplier onboarding.
Legacy system integration	Existing ERP, WMS, and quality systems were not designed to emit or consume standardized event data.	Deploy middleware/API layers rather than full system replacement; phase integration by the highest-risk product line first.
Cost and unclear ROI	Sensor deployment, connectivity, and platform costs are visible and immediate; risk-reduction benefits are probabilistic and delayed.	Pilot on highest-risk SKUs; quantify recall-cost avoidance and insurance/audit benefits explicitly
Data quality and sensor reliability	Sensor drift, gaps, and manual data-entry errors undermine model accuracy and trust in outputs.	Implement automated data-quality checks and sensor calibration schedules; flag low-confidence records rather than silently dropping them.
Multi-party adoption and coordination	Full-chain traceability requires every partner to participate; a single non-compliant link breaks the chain.	Use phased, risk-prioritized rollout; provide low-cost onboarding tools for smaller suppliers.
Data privacy and competitive sensitivity	Partners are reluctant to share commercially sensitive volume, pricing, or sourcing data.	Apply role-based access controls and data-minimization principles; share event metadata rather than full commercial detail.

11. Industry Applications and Emerging Practices

For the consumer products and food manufacturing industry, traceability investment has focused initially on high-risk, high-regulatory-scrutiny products, such as fresh produce, seafood, infant formula and/or other high-risk products listed on regulatory traceability lists, and high-value proteins. Among the first and most obvious examples of blockchain-powered EPCIS-based traceability is the case of fresh produce supply chains, which have found some contamination scandals in the leafy-green market to be a major motivator, partly because, of all the contamination cases that happened in the produce sector, they found paper-based trace-backs to be slow at scale. Combining DNA-based verification and digital chain-of-custody documentation has been implemented in seafood supply chains, where the risks of species substitution fraud and illegal, unreported, and unregulated fishing are exceptionally high, to address both food safety and chain provenance at the same time [21].

In addition to being able to trace back for food safety purposes, packaged and shelf-stable consumer product manufacturers are expanding their fresh and perishable traceability systems to other parts of the supply chain, such as tracking suppliers' financial health, verifying the ethics behind sourcing, or tracking carbon footprints—systems that can be built around the same underlying event infrastructure, with comparatively small incremental investments.

Contract manufacturers and private label manufacturers that may have varying contract requirements for product traceability can have unique issues when considering adoption, as having a separate data pipeline for every customer relationship is generally not cost-effective. Many have responded to this by creating a single internal EPCIS-compliant event infrastructure on top of which they then expose views or extracts of their customer-specific relationships—effectively using the element of traceability data capture as a shared operational tool instead of a brand-specific deliverable for any one of their relationships—and this approach is being increasingly replicated by supply chain participants other than customers, such as third-party logistics providers or co-packers.

One common theme that emerges from these early efforts is that successful programs begin very small, on a single high-risk product category or a single critical tracking event – for example, receiving or shipping—and grow horizontally over time, having proven the worthwhile data model, partner onboarding process, and governance.

12. Future Directions

There are several new capabilities to be expected in the next few years as part of the traceability architecture outlined in this article. Continuously feeding Data Record Types (DTRs) by the same event stream from EPCIS will enable organizations to anticipate the consequences of disruptive or contaminative events in the supply chain that are linked upstream to them by simulating the impact of their downstream events before they actually happen, thereby aiding proactive risk mitigation, rather than only reactive response. It is being used on the unstructured side of traceability data, automatically summarizing supplier audit reports and certificate of analysis documents and drafting regulatory notification language during an active incident and answering natural language trace-back requests from quality and compliance personnel without the need to be an expert in query language.

These should be considered as workforce and organization-relevant technical shifts. The skill set of quality and food safety teams that, until quite recently, have worked mainly as manual inspectors and paper-trail investigators but now must employ the ability to interpret model outputs, estimate confidence intervals, and recognize when a "red flag" is there due to a confidence interval that actually points to a green light rather than a red flag is a skill set that has not been widely taught in food science and food safety programs. Companies that invest in cross-training quality personnel alongside their data science and IT teams fare better in getting data science tools and artificial intelligence/machine learning into the day-to-day business decision-making process and doing so more swiftly and sustainably—not delegating capability to a dedicated analytics team.

Predictive recall capability is an additional step that builds on risk-scoring techniques from Section 5, where models continuously use leading indicators, including audit outcomes, weather events at growing regions, and trends in past incidents, to determine the likelihood that a quality event has occurred at a supplier or facility, potentially enabling earlier intervention to prevent contaminated product from entering the supply chain, rather than simply putting the brakes on when it happens. Finally, the evolution of standards for interoperability and shared supplies of industry data trusts is expected to reduce the costs of multi-party involvement, especially for smaller suppliers, and is a top remaining obstacle to achieving truly broad end-to-end coverage.

13. Conclusion

In the food safety industry, end-to-end supply chain visibility and traceability have progressed from a compliance requirement to a strategic asset that significantly impacts food safety results, operational resilience, and consumer confidence in the consumer product sector. In this article, the ravenous reference architecture, AI/ML use cases, and incident workflow presented are a clear and attainable journey from raw data from data sensors and data from incidents to data from active risk intelligence and, finally, to faster and more accurate recall and response. Achieving this potential is more of a multi-party project than it is about any single technology and instead relies on the sustained collaboration of multiple parties throughout a supply chain's myriad independent players, as well as the adoption of standards and disciplined data governance. Organizations that see traceability as an operational system that is continually improved and that leverages the power of AI will be best placed to transform the need for compliance into a long-term competitive advantage.

References

[1] Guillaume Jean, “Supply Chain Visibility and Transparency,” 2024. [Google Scholar]

[2] Amy Schmidt, What is End-to-End Traceability?, 2023. [Online]. Available: https://producersmarket.com/blog/what-is-end-to-end-traceability/

[3] Ankit Singh, Balancing Food Safety and Sustainability in Food Processing Systems, 2026. [Online]. Available: https://www.azocleantech.com/article.aspx?ArticleID=2093

[4] Richard Wilson, FSMA 204 Timeline Shifted: Key Takeaways for Food Businesses, Nulogy, 2025. [Online]. Available: https://nulogy.com/fsma-204-timeline-shifted-key-takeaways-for-food-businesses/

[5] Food Safety Magazine, FDA Traceability Rule Readiness Exercises Reveal Supply Chain Coordination Matters More Than Technology, 2026. [Online]. Available: https://www.food-safety.com/articles/11507-fda-traceability-rule-readiness-exercises-reveal-supply-chain-coordination-matters-more-than-technology

[6] Navigating Global Food Regulations: Safety, Labelling and Compliance, 2026. [Online]. Available: https://www.sial-network.com/en/trends/Newsroom/reglementations-alimentaires-mondiales-securite-etiquetage-conformite

[7] N. Yerashevich, EU Traceability and Transparency Requirements: Why They Now Matter Beyond Agri-Food, Ohana Public Affairs Consultancy, 2026. [Online]. Available: https://ohanapublicaffairs.eu/2026/05/06/why-eu-transparency-and-traceability-requirements-now-matter-far-beyond-agri-food/

[8] Sylvain Charlebois et al., "Comparison of Global Food Traceability Regulations and Requirements," *Comprehensive Reviews in Food Science and Food Safety*, vol. 13, no. 5, pp. 1104-1123, 2014. [CrossRef] [Google Scholar] [Publisher Link]

[9] GSI, GS1 Global Traceability Standard, GSI, 2017. [Online]. Available: https://www.gs1.org/standards/gs1-global-traceability-standard/current-standard

[10] Amogh Balaji, End-To-End Traceability: Connecting Every Critical Link in the Global Supply Chain, The Traceability Hub, 2026. [Online]. Available: https://thetraceabilityhub.com/end-to-end-traceability-global-supply-chain/

[11] GS1, GS1 Traceability Reference Book, 2021. [Online]. Available: https://www.gs1.at/sites/default/files/2022-02/gs1-traceability-reference-book-2021.pdf

[12] Cargoz, One Up/One Down. [Online]. Available: https://www.cargoz.com/glossary/one-up-one-down-1375

[13] Rodolfo Bojorque et al., "A Comparative Study of Unsupervised Machine Learning and Deep Learning Techniques for Anomaly Detection in Recommender Systems," *Information*, vol. 17, no. 5, pp. 1-23, 2026. [CrossRef] [Google Scholar] [Publisher Link]

[14] Rahima Khanam et al., "A Comprehensive Review of Convolutional Neural Networks for Defect Detection in Industrial Applications," *IEEE Access*, vol. 12, pp. 94250-94295, 2024. [CrossRef] [Google Scholar] [Publisher Link]

[15] Contextual Computing, Traceable Machine Learning Governance: Standards, Organizational Initiatives and Automation, 2018. [Online]. Available: https://contextualcomputing.net/2018/01/29/traceable-machine-learning-governance-standards-organizational-initiatives-and-automation/

[16] Emily Winks, How to Version AI Models for Reproducibility and Governance, Atlan, 2026. [Online]. Available: https://atlan.com/know/ai-model-versioning-best-practices/

[17] Artem Storozhuk, and Anastasiia Voitova, Audit Logs Security: Cryptographically Signed Tamper-Proof Logs, Cossack Labs, 2020. [Online]. Available: https://www.cossacklabs.com/blog/audit-logs-security/

[18] R. Hanumantharaju et al., "Blockchain Based Machine Learning Approach for Secure and Efficient Vehicular Data Monitoring and Analysis," *Discover Computing*, vol. 28, no. 1, 2025. [CrossRef] [Google Scholar] [Publisher Link]

[19] Konstantinos Rotsios et al., "Evaluating the Use of QR Codes on Food Products," *Sustainability*, vol. 14, no. 8, pp. 1-15, 2022. [CrossRef] [Google Scholar] [Publisher Link]

[20] Loftware, From QR Codes to Connected Experiences: The New Era of Food & Beverage Packaging. [Online]. Available: https://www.loftware.com/resources/white-papers/2025/from-qr-codes-to-connected-experiences-the-new-era-of-food-and-beverage-packaging

[21] Anthony J. Silva, and Rosales S. Hellberg, "DNA-based Techniques for Seafood Species Authentication," *Advances in Food and Nutrition Research*, vol. 95, pp. 207-255, 2021. [CrossRef] [Google Scholar] [Publisher Link]