June 2026

Interoperable FSMA 204 Traceability

Lessons from the PFT Cucumber Pilot

Partnership for Food Traceability

Executive Summary

FSMA 204, the FDA's Food Traceability Rule, requires companies that handle foods on the Food Traceability List to capture standardized Key Data Elements at defined Critical Tracking Events and to produce an electronically sortable record of that data within 24 hours of a request. For most organizations, the challenge is not capturing data within their own four walls but moving it reliably across a fragmented supply chain of growers, packers, processors, distributors, and retailers, each running different systems. Compliance at scale, therefore, depends on interoperability rather than on any single platform.

To test whether interoperability is achievable in practice, the Partnership for Food Traceability convened a pilot in which three solution providers, Kezzler, Peer Ledger, and Starfish, simulated an FSMA 204 traceability system for a cucumber supply chain running from farm through packing, slicing, cold storage, and deli salad manufacturing to distribution and retail. Each provider owned a different portion of the chain, GS1 identifiers and EPCIS event structures carried the data, and APIs moved events between systems in near real time.

The pilot confirmed that multiple independent platforms can exchange GS1-aligned traceability data and assemble an end-to-end, electronically sortable record set that meets the structural expectations of the rule. It also showed that the hardest work is not the API connections but the surrounding agreements: aligning master data, deciding when and where traceability lot codes are created, translating between internal event formats and the flattened output FDA expects, and defining who may see which data under what circumstances.

The central recommendation is to treat FSMA 204 as a network problem rather than a company-by-company one. Practical steps include designating a single party to own canonical master data, agreeing on lot code conventions early and writing them into contracts and configuration, treating the translation from operational data to regulatory output as a deliberate design task, and investing in structured onboarding and clear permissioning before relying on these connections during an actual recall. Interoperable traceability is achievable today with existing standards, and much of the practical value

lives in the interoperability layer that normalizes, routes, and reconciles data across platforms.

Introduction

The U.S. Food and Drug Administration's Food Traceability Rule, which implements Section 204 of the Food Safety Modernization Act (FSMA 204), requires companies to maintain standardized "Key Data Elements" (KDEs) at defined "Critical Tracking Events" (CTEs) and to provide those records to FDA in an electronically sortable format within 24 hours of a request.

While the regulatory expectations are now clear on paper, the practical question for industry is how to exchange this data across a fragmented supply chain of growers, packers, processors, distributors, and retailers, each using different technology providers and internal systems. Achieving this in a way that is repeatable, scalable, and economically sustainable requires an interoperable approach to data sharing.

The Partnership for Food Traceability (PFT) was created as a public-private, sector-neutral collaboration to help answer these questions. PFT's mission is to accelerate practical, industry-led solutions that can help companies move from interpreting FSMA 204 requirements to implementing them in real supply chains.

This whitepaper describes a pilot data-sharing project hosted by PFT in which three technology solution providers, Kezzler, Peer Ledger, and Starfish, worked together to simulate an interoperable FSMA 204–ready traceability system for cucumbers. The simulation follows product and data from farm through initial packing and processing, into distribution, and finally to retail. Using GS1-compliant identifiers and data structures, the pilot focuses on how independent systems can collaborate to deliver a single, coherent traceability view to regulators and trading partners.

The paper is intended for supply chain, food safety, IT, and compliance leaders who are responsible for implementing FSMA 204 in complex, multi-partner environments. It documents the pilot design, summarizes the technical approach, and highlights key successes, challenges, and lessons learned that can inform production-grade implementations.

FSMA 204 and the Interoperability Challenge

The FSMA Final Rule on Requirements for Additional Traceability Records for Certain Foods (colloquially referred to as FSMA 204) requires organizations that manufacture, process, pack, or hold foods on the Food Traceability List (FTL) to retain and share KDEs linked to a traceability lot code across a series of CTEs, including growing, receiving, transforming, creating, and shipping. For many companies, the main challenge is not capturing data within their own four walls, but ensuring that data can move with the product as it crosses organizational and system boundaries.

Today's food supply chains are characterized by heterogeneous data collection systems, including on-farm tools, pack-house platforms, ERPs, warehouse management systems, and specialized traceability solutions. Different parties often use different identifiers, data models, and integration patterns. GS1 standards and EPCIS provide an important common language, but implementing them consistently across multiple providers is non-trivial.

The cucumbers pilot was designed against this backdrop. It assumes that no single solution provider will ever own the entire stack, and that FSMA 204 compliance will depend on interoperable networks of platforms. The pilot's central question is therefore not "Can one platform comply?" but "Can multiple platforms, each playing a different role in the supply chain, exchange data in a way that meets FSMA 204 expectations and supports rapid, accurate trace-back and trace-forward?"

Overview of the PFT Cucumber Pilot

The pilot uses a stylized but realistic cucumber supply chain as a testbed (Figure 1). Product originates at multiple cucumber farms, moves through initial packing, is transformed at a processor into consumer-facing products, and then flows through distribution centers to retail stores where consumers purchase the product. Each step in this chain represents one or more FSMA 204 CTEs and requires specific KDEs to be captured and linked via a traceability lot code.

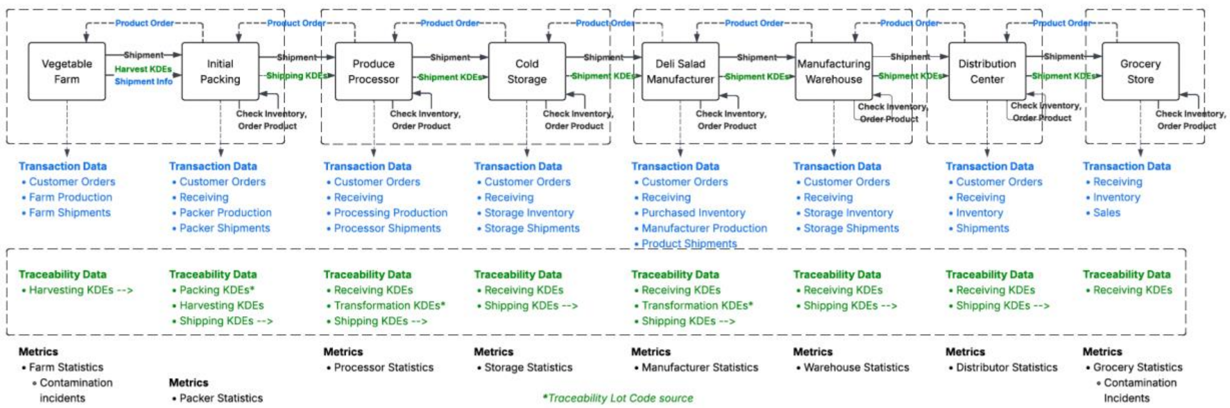

Figure 1. Simplified Supply Chain Diagram used in the Partnership for Food Traceability (PFT) Cucumber Interoperability Pilot. Peer Ledger handled CTEs for the farm, packer, slicing processor, cold storage, and distribution center, while Kezzler handled CTE data for the deli salad processor, manufacturing warehouse, and retailer nodes. Starfish served as the central data harmonization and interoperability layer for the pilot. Starfish's platform ingested, normalized, and unified traceability data from both Kezzler and Peer Ledger for all companies in the pilot.

PFT convened three solution providers to participate in the pilot:

Kezzler provides a highly scalable, enterprise-grade traceability platform built to track products across global supply chains. For this pilot, Kezzler acted as the traceability layer for the *processor* and *retailer*, establishing a continuous digital thread throughout the product journey.

At the processing stage, the platform managed the transformation of raw cucumbers into a deli salad, maintaining full data cohesion while assigning a new Traceability Lot Code (TLC). The system tracked the product through its entire lifecycle, including intracompany transfers to its manufacturing warehouse and external dispatch to the distribution center. As the grocery store operator, Kezzler successfully integrated incoming shipping KDEs from the distribution center, demonstrating the platform’s ability to reconstruct the complete flow from raw ingredient receipt through processing to final retail sale. This creates the potential for bidirectional data sharing, providing upstream owners with downstream visibility into their product's journey.

The platform accomplishes this by harmonizing data from diverse sources, structuring events in GS1 EPCIS format, and associating data points (including FSMA 204 KDEs) with digital objects such as lots or batches. This approach ensures that structured data is always ready for on-demand FDA reporting while providing stakeholders with clear, real-time visibility through dashboards and granular data filtering.

Peer Ledger gives enterprises a patented, audit-ready platform to trace every food product, component, and raw material across their global supply chain and to know it's continuously verified, in near real-time. Peer Ledger has a highly scalable, blockchain- and EPCIS-based, NIST IR 8536-compliant traceability platform that supports both GS1 and non-GS1 labeled products. Peer Ledger's software collects bilaterally verified, reliable, and up-to-date supply chain traceability data on a continuous basis, assuring data accuracy.

In the pilot, Peer Ledger represented the traceability layer for the farm, packer, slicing processor, and cold storage in the upstream and midstream. Its single downstream node was the distribution center. Peer Ledger delivered a continuous, verified digital thread connecting the upstream and midstream supply chain participants, and demonstrated that its platform interoperates fully with Kezzler via Starfish to support the FDA's requirement for end-to-end product traceability across the whole supply chain.

Peer Ledger recorded the appropriate CTEs and tracked the cucumber product from farm to packer, packer to first processor (slicer), and slicer to cold storage and then to the deli salad manufacturer. At the cucumber farms, a harvest CTE with shipping information for bulk cucumbers was recorded on Peer Ledger's system. The bulk cucumbers were then received at the packers and put into cases. At the packer, Peer Ledger made the first assignment of a Traceability Lot Code (TLC) to a cucumber product in this supply chain. The TLC was made up of a GTIN+LOT CODE. Every case of cucumbers from the same lot had the same TLC.

The packer shipped the cases of cucumbers to the slicer processor. At a slicer, the cucumbers were sliced (transformation event was recorded) and put into new cases with new TLCs. A shipping event was recorded for sending the cases of sliced cucumbers to cold storage. The shipping event from the cold storage to the deli cucumber salad manufacturer triggered interop handoffs to and from Starfish for unification with Kezzler's tracing. Kezzler, as the grocery store operator, successfully integrated incoming shipping KDEs via Starfish from the distribution center operated by Peer Ledger, demonstrating that every step of the product life cycle, from raw ingredient receipt through processing to point of sale, can be fully reconstructed, verified, and audited on demand.

Starfish served as the central data harmonization and interoperability layer for the pilot. Starfish's platform ingested, normalized, and unified traceability data from both Kezzler and Peer Ledger for all companies in the pilot, reconciling differences in data formats (EPCIS XML, FSMA reports) and identifier conventions across systems, determining access controls for what data should be shared with each individual pilot company, and feeding the relevant data back into Kezzler and Peer Ledger for them to present to each company. In the supply chain simulation, Starfish represented the connective tissue between the systems enabling each technology system to seamlessly facilitate data sharing as required by FSMA 204. To be clear about the division of labor, Starfish operated as the interoperability layer between systems, normalizing and routing data and reconciling identifiers rather than serving as a single system of record. Each provider's platform assembled the events it received into its own traceability view.

The pilot uses GS1 identifiers (for example, GTINs and GLNs) and ASNs based on EDI 856 to represent CTEs such as harvest, pack, ship, receive, transform, and sell. Events are exchanged via APIs rather than batch EDI to reflect the emerging preference for near real-time traceability data sharing.

The project is explicitly experimental, and its purpose is to expose friction, test different integration patterns, and document what it actually takes for independent solutions to interoperate in ways that are aligned with the text and intent of FSMA 204.

Pilot Design: From Cucumber Farm to Checkout

In the simulation and consistent with the Rule, cucumbers are assigned traceability lot codes during the initial packing CTE, and those codes are carried through subsequent steps: initial shipping, receiving into processing, transformation into consumer units, case-level and pallet-level aggregation, distribution movements, and retail sale.

Each solution provider had a defined role and vantage point in the simulated chain. For example, when Peer Ledger generated farm-level harvest and packing events for a cucumber lot, that data was fetched by Starfish, normalized, and then automatically pushed to Kezzler's EPCIS capture endpoint. Kezzler's platform unified these initial records with subsequent transformation and downstream events, creating a continuous traceability thread within its scalable EPCIS repository. As Kezzler recorded these later events, Starfish retrieved the updated data to provide Peer Ledger with relevant shipping notifications. This bidirectional, hub-and-spoke pattern allowed each provider to work within its native data format while Starfish handled translation and routing.

The pilot design deliberately introduces complexity:

- Multiple farms contributing to a single processing batch.
- Multiple SKUs and packaging hierarchies created from common incoming lots.
- Shipping events that aggregate cases and pallets.

These features are essential to stress-test how well interoperable systems can maintain an unbroken digital thread and how easily they can generate an FDA-ready, electronically sortable spreadsheet when given a suspected lot or retail unit as a starting point.

Execution and Data Flows

The pilot was executed as an iterative simulation rather than a one-time data load. Participants worked through a series of design sessions and working meetings to agree on:

- How to express master data for farms, facilities, and products in a way that each system could consume.
- How to encode CTE events and map them to FSMA 204 KDE requirements.
- How to move event data between systems via APIs while preserving key relationships such as aggregation and transformation.

From there, each provider generated test data for their portion of the chain, and the team exercised different data-sharing patterns.

A key outcome of execution was the ability to assemble end-to-end traceability views across the three platforms and to generate an electronically sortable record set that shows, for a given retail unit or lot, the complete upstream path and associated CTEs.

Key Findings from the Pilot

At a high level, the pilot demonstrated that multiple independent solution providers can, in fact, exchange GS1-aligned traceability data in a way that satisfies the structural expectations of FSMA 204. It also showed that ASNs and modern APIs can significantly reduce the friction, compared to paper-based or manual data capture systems.

However, the exercise also surfaced important nuances that are easy to overlook when discussing FSMA 204 at a conceptual level. Master data assumptions, timing of lot identifier assignment, subtle differences between input and output schemas, and the mechanics of permissioning all had material impacts on how smoothly the systems could interoperate. These nuances form the core of the “Lessons Learned” and “What Worked Well” sections below.

Lessons Learned

Master Data - Lot Identifiers Come Earlier Than Expected

One of the clearest lessons from the pilot is that traceability lot identifiers could appear earlier in the process than many participants initially expect. Systems must anticipate this edge case. In the cucumber scenario, traceability lot codes are effectively “born” at the pack-house versus any codes used at a farm and must be consistently referenced through packing, processing, distribution, and retail.

This has implications for how master data is structured and shared. If upstream parties introduce lot codes that downstream systems do not anticipate, or if multiple identifiers are used for what is, in practice, the same logical lot, it creates ambiguity that undermines traceability. The pilot emphasized the need for an early, shared agreement on how lots will be defined, when they are created, and how they will be referenced across platforms.

In a production environment, this likely means that master data stewardship cannot be treated as a purely local concern. Suppliers, processors, and technology providers need explicit conventions for lot definition, naming, and lifecycle, and these conventions should be encoded in both contracts and configuration.

The pilot also revealed that master data such as organizations, locations, products, and their associated GS1 identifiers must be treated as a controlled, centralized input rather than allowing each party to independently define these entities. In the early phases of the pilot, master data was maintained in a shared spreadsheet, which led to version drift as different parties worked from slightly different copies. Small discrepancies in location identifiers, for example, were difficult to detect until actual event data began flowing through the systems and failed to reconcile. The team ultimately adopted an approach where one party (Starfish) generated the canonical set of organizations, GTINs, and GLNs, and other participants referenced those identifiers when generating their event data. This reduced ambiguity significantly and is a pattern we would recommend for any multi-party traceability implementation.

Data Translation - Inputs Are Different Than Outputs

The pilot also highlighted that, in practice, the data structures required for internal processing and those required for FSMA 204 reporting are not identical. A concrete example is shipping events. Within an EPCIS-style event repository, a shipping event may be represented as a single event containing multiple line items or aggregated objects. That structure works well for operations and system-to-system messaging.

However, when generating the electronically sortable spreadsheet that FDA expects under FSMA 204, the team found it necessary to “flatten” these rich events into multiple rows, effectively splitting a single shipping event with multiple items into individual rows for each product, lot, or aggregation unit. This translation step is not complicated, but it is essential.

The lesson is that organizations should not expect their operational event model and their regulatory reporting format to be one-to-one. Instead, they should plan for a translation layer that can take EPCIS-style input and produce spreadsheets that conform to FSMA 204 expectations without losing important context about aggregation and transformation relationships.

From Starfish's perspective as the normalization layer, a key learning was the importance of pre-built, standardized connectors for common data formats (EPCIS 1.0, EPCIS 2.0, FSMA-specific schemas, EDI). At the outset of the pilot, ingestion logic was built ad hoc for each partner, creating a direct pipeline from each partner's format into Starfish's internal model. Over the course of the pilot, this was refactored into a three-stage architecture: partner data is received in its native format, normalized to the partner's logical model, and then converted to Starfish's internal model. This approach made it far easier to identify where data discrepancies originated and will reduce onboarding time for future partners.

From Kezzler’s perspective, the pilot showed that while operational events and regulatory reports use different formats, the information behind them should be the same. By linking event data to a digital ID like the Traceability Lot Code (TLC), raw data becomes a reusable record of the product journey. This makes it possible to generate the specific reports the FDA requires while keeping the data flexible for other needs. Ultimately, this ensures the information is ready for compliance, sharing with partners, and internal analysis.

Connectivity and Permissioning - Getting Access Right

From a connectivity standpoint, the pilot underscored how central permissioning is to any interoperable traceability network. Even in a simulated environment, decisions had to be made about who could see what data, under what circumstances, and with what level of granularity.

Building API connections between systems was relatively straightforward compared to designing a permissioning model that balanced privacy, commercial sensitivity, and data access in the event of an incident. Role-based access, data sharing agreements, and thoughtful use of technologies are all tools that can help, but they require deliberate design rather than being bolted on at the end.

In operational deployments, this will likely mean that trading partners must agree up front on governance. Agreement is needed about when data is shared automatically, when explicit consent is required, and how exceptional circumstances (for example, regulator-requested investigations) trigger broader access.

From Starfish's vantage point as the routing layer, permissioning was implemented through a "Connector Organization Document" model, where each organization in the supply chain was explicitly associated with the connector (and therefore the external system) that should receive its events. This meant that when a new event was ingested for a given organization, the push service knew exactly where to route it without requiring broad data sharing. This approach kept data flows narrow and intentional, but it also required careful setup as the association between organizations and connectors had to be established after master data ingestion, and any misconfiguration would result in events silently failing to propagate. In production, this type of routing configuration should be part of a formal onboarding checklist.

What Worked Well

Simulation - Building Out a Realistic Scenario

A major success factor for the pilot was the decision to treat the exercise as a realistic simulation rather than a purely technical demonstration. The cucumber supply chain included multiple farms, a pack-house, two processors (slicer and deli salad manufacturer), distribution centers, and retail outlets, with plausible product hierarchies and movement patterns.

This level of realism provided immediate value. It allowed participants to test assumptions about where CTEs actually occur, how KDEs are captured at each step, and how data must move when lots are split, merged, or re-packed. It also made the resulting lessons more relevant to industry stakeholders who will recognize similar patterns in their own operations.

The simulation was also refined iteratively over several months, with the dataset and data flow conventions evolving as the team encountered real integration friction. This iterative approach, rather than a "big bang" data load, was essential to surfacing issues like format mismatches and identifier inconsistencies early enough to resolve them.

Interoperability - Shared Connections Enable a More Efficient Flow of Data

The pilot confirmed that when solution providers align around common identifiers (such as GTIN, GLN, and traceability lot codes) and shared event semantics, data can flow efficiently between platforms without extensive custom mapping for each trading partner.

As long as each platform could create and consume CTE structures, the integration overhead for adding a new participant to the network remained manageable.

Coordination - Regular Cadence Helps Maintain Momentum

Another practical success was the decision to maintain a regular meeting cadence across all participating organizations. The pilot involved not only three solution providers but also PFT coordination, and subject matter experts in FSMA 204. This group met weekly for several months from August to December 2025. Without structured coordination, it would have been easy for the work to fragment into isolated technical efforts.

Regular working sessions allowed the group to resolve questions about data standards, align on event semantics, and make decisions about scenario changes. They also kept the pilot moving forward despite competing priorities at each participating organization.

The experience suggests that any multi-party traceability initiative, whether in produce or another FTL category, should treat governance and coordination as first-class workstreams alongside technology and data modeling.

Cost Efficiency - API-Based and Connector-Driven

While the pilot did not produce formal cost benchmarks, one notable observation is that the API-based, connector-driven architecture allowed integration work to proceed with relatively small engineering teams at each organization. The bulk of the technical effort was concentrated in initial connector development and master data alignment; once those were in place, adding new event types or supply chain nodes required incremental rather than ground-up work. This suggests that the marginal cost of extending an interoperable traceability network decreases as shared conventions and connectors mature.

Performance - APIs Are Fast and Efficient

The pilot also demonstrated that modern API-based architectures can support the performance demands of traceability queries. Even with multiple systems participating and a non-trivial number of events, queries for a given lot or retail unit could be answered promptly.

While performance will always depend on implementation details and infrastructure choices, the pilot provides early evidence that FSMA 204 traceability does not have to come at the expense of query speed, especially when systems are designed from the outset with event-driven APIs and efficient indexing.

Implications for Industry

The cucumber pilot offers several takeaways for companies that are designing their own FSMA 204 programs. First, interoperability is achievable today using existing standards and technologies, and does not require a single, monolithic platform. Second, the most challenging work is often not the APIs themselves but the alignment on master data, event semantics, and permissioning. Third, simulation exercises can be a powerful way to test assumptions before rolling out large-scale implementations.

For many organizations, the practical implication is that FSMA 204 efforts should be organized around networks rather than individual companies. Processors, distributors, and retailers will need to sit down with their technology partners and upstream suppliers to agree on shared conventions and to test those conventions in controlled pilots before relying on them in real incidents.

Recommendations and Next Steps

- Encourage multi-party simulations that mirror real supply chains before committing to large-scale roll-outs.

- Formalize master data and lot creation rules early, and reflect them in contracts and configuration across all platforms.
- Designate a single party as the master data steward for each pilot or implementation, responsible for generating and distributing canonical identifiers (GTINs, GLNs, organization records) to avoid version drift across participants.
- Invest in a well-defined onboarding flow for new technology partners, including a standard set of APIs, expected data formats, and validation checks. The PFT pilot demonstrated that even GS1- or EPCIS-compatible systems exhibit subtle differences in data representation that are best caught through a structured onboarding process rather than discovered during live data exchange.
- Treat the translation from operational event data to FSMA 204-ready spreadsheets as a distinct design problem, not an afterthought.
- Invest in clear governance and permissioning frameworks so that APIs and events operate within an agreed trust boundary.

Future phases of the PFT work may extend the pilot to additional FTL commodities, introduce more trading partners and solution providers, and test more complex scenarios such as rework, returns, or cross-docking. Future pilots may also explore challenges dealing with large datasets, as event data can be modeled to scale over a time period of several years.

Conclusion

FSMA 204 is accelerating the food industry's shift from fragmented, paper-heavy records to connected, digital traceability networks. The PFT cucumber pilot shows that when solution providers align on standards, share a clear scenario, and invest in coordination, interoperable data sharing is not just a theoretical goal.

The pilot reinforced that the emerging traceability ecosystem will be a network of specialized platforms rather than a single monolithic solution, and that the interoperability layer, defined as the ability to normalize, route, and reconcile data across those platforms, is where much of the practical value and complexity resides.

By documenting the pilot's design, what worked, and where friction emerged, this whitepaper aims to provide a roadmap for companies that want to move beyond minimum compliance and build traceability systems that support faster investigations, more efficient operations, and greater transparency across the supply chain.

About the Authors

Benjamin Miller, Ph.D., MPH

COO & EVP of Regulatory and Scientific Affairs, The Acheson Group (TAG)

Dr. Miller has over twenty-five years of experience in food safety, public health, and regulatory affairs. He is a founding member of the FDA-sponsored Partnership for Food Traceability (PFT) and an expert in FDA traceability requirements, having previously led the State of Minnesota's food regulatory programs as Division Director at the Minnesota Department of Agriculture.

Wiggs Civitillo

Founder & CEO, Starfish

Wiggs Civitillo leads Starfish, a supply chain connectivity platform enabling secure, standards-based data exchange across the food ecosystem. He has spent his career advancing trusted interoperability to help build safer, more resilient supply chains, including leading the global food traceability practice at IBM and working with industry groups like PFT, GS1, GDST, SCOTF, and IFMA. He speaks frequently on FSMA 204, food traceability, and the practical work of connecting fragmented supply chains.

Frederic Clulow

General Manager, Kezzler Americas

Frederic Clulow brings two decades of expertise in supply chain traceability across industries including pharmaceuticals, semiconductors, and food. He leads Kezzler's commercial development in the Americas and is a recognized voice on FSMA 204 implementation, digitization of traceability, and the importance of interoperability across the food supply chain.

Dawn Jutla, Ph.D.

Founder & CEO, Peer Ledger Inc.

Dr. Jutla is the founder and CEO of Peer Ledger, whose patented, trusted blockchain-based traceability platform can track any GS1, non-GS1, or GDST product, its components and packaging in the food or other industries. She holds a Ph.D. in Computer Science and has served as a Full Professor at the Sobey School of Business for the past decade. Her academic and industry work spans research in distributed systems, electronic commerce, and supply chain trust.

Bob Celeste

Founder, Center for Supply Chain Studies

Bob Celeste has over thirty years of experience in standards, technologies, and supply chain operations across pharmaceuticals, consumer goods, fresh foods, and other industries. Prior to founding the Center, he served as Senior Director at GS1 US, where he led standards-based, item-level traceability initiatives. He frequently advises the FDA, state regulators, and international governing organizations on traceability implementation.