SUPPLY CHAIN RESEARCH · EDITORIAL FEATURE

The Traceability Mandate

Why Item-Level Compliance Is Bought as Software and Fails as Data Supply

Three item-level traceability regimes set statutory deadlines. All three slipped by years, and the stated reasons converge on the same point: the software was ready and the data was not. One grocery retailer moved its supply base faster than any of them.

Independent, vendor-neutral research for supply chain and technology leaders
Supply Chain Research | supplychainresearch.com | 2026

Contents

Why Item-Level Compliance Is Bought as Software and Fails as Data Supply

01 Executive summary

Item-level traceability arrived as a regulatory obligation and was received as a procurement exercise. Pharmaceutical serialisation, food traceability, deforestation-free sourcing, and forced-labour import controls each require an organization to know, and to be able to demonstrate on demand, where a specific unit came from and every hand it passed through. Each obligation has produced a market of platforms that capture events, serialise units, exchange records between trading partners, and generate the reporting a regulator would want. Those platforms broadly work. The programmes built on them have nonetheless slipped by years, and the reason is consistent across regimes: the software was ready and the data was not.

The evidence for that claim is unusually direct, because the regulators themselves have said so. The pharmaceutical deadline of November 2023 was followed by a one-year stabilization period and then by staggered exemptions running to November 2026, granted by trading-partner class in descending order of size and capability. The food traceability rule was extended thirty months, from January 2026 to July 2028, and the regulator's stated reason was that even prepared entities depend on receiving accurate data from supply chain partners who may not be similarly situated. The deforestation regulation was delayed twice, most recently because the information system could not handle the projected volume of due-diligence statements. We say honestly that these mandates are worth having, that the public-health case behind the food rule is documented in outbreak investigations where a slow traceback cost lives, and that item-level traceability has delivered real benefit where it has been implemented. The argument here is narrower: the binding constraint is supplier capability and the leverage to compel it, not software, and a programme scoped as a technology purchase is scoping the wrong problem.

30 months
the extension granted to the food traceability rule, from January 2026 to July 2028

4
distinct trading-partner classes granted separate pharmaceutical exemption dates, staggered by capability

Aug 2025
the month one grocery retailer's own traceability mandate took effect, three years ahead of the federal rule

The argument in brief

- **The platform is the last ten percent.** Event capture, serialisation, and exchange are solved problems. Whether complete, standardised, accurate data arrives from every participant is not, and it sits outside the buyer's systems.
- **Every regime slipped, and for the same reason.** Pharmaceutical, food, and deforestation deadlines each moved by years. The stated justifications converge on readiness of the data supply rather than availability of technology.
- **The staggering is the evidence.** Exemptions granted by trading-partner class in descending order of size are a regulator conceding that capability, not compliance intent, is the constraint.

- **Leverage moves supply bases; deadlines do not.** One grocery retailer required traceability event data with financial consequences from August 2025 and its supply base moved, while the federal rule covering similar data slipped to 2028.
- **Almost every benefit figure is sold by someone.** Recall-time savings and cost-avoidance claims originate with compliance vendors. The regulator-authored outbreak research is the evidence a business case should rest on.

The pattern this article describes is not confined to the four regimes examined. Any obligation that requires unit-level identity to survive multiple changes of custody, and that places the compliance burden on a party downstream of most of the events, will encounter the same constraint. That description fits several regimes now in preparation, which means the lesson has forward value rather than being a retrospective on programmes already underway.

02 Bought as software, failing as data supply

The structural claim can be stated in a sentence and is worth stating before the regulatory detail, because the detail is easier to follow once the shape is clear. A traceability obligation is a requirement about an entire chain of custody, and it is discharged through a procurement made by a single organization at one point in that chain. Figure 1 sets out the mismatch.

Figure 1. The mandate is bought as software and fails as data supply. The platform is budgeted, scoped, and demonstrable. Whether complete and accurate event data arrives from every participant is outside the buyer's system and largely outside its control.

Consider how such a programme is typically initiated. A compliance obligation is identified, usually by legal or regulatory affairs. It is translated into a set of system requirements, because the obligation is expressed in terms of records that must be captured, retained, and produced. Those requirements go to a technology function, which runs a selection process and procures a platform. The platform is implemented, integrated with the enterprise systems, and tested. At the end of this process the organization has a capable system, a defensible audit trail of its own compliance effort, and a programme that reports as substantially complete.

What the organization does not have, and what the obligation actually requires, is a complete record of events that occurred at parties it does not own. Traceability regimes are constructed around critical tracking events and the key data elements associated with them, and those events occur at every point where a product changes hands or changes state: harvested, cooled, packed, shipped, received, transformed. Most of those events occur before the product reaches the organization that bought the platform. The record can only be complete if every upstream participant captures its events in a compatible form and transmits them, and the organization's platform cannot make that happen.

The asymmetry is what makes this failure mode so persistent. The part of the problem the organization controls is the part it can budget for, scope, implement, and demonstrate, so that part gets done and gets reported. The part that determines whether the obligation is actually met sits with hundreds or thousands of upstream parties whose capability varies enormously, who have no contract with the organization beyond the first tier, and who face costs to comply that no one has offered to fund. A programme that measures its progress by platform implementation will report green throughout, right up until a regulator asks for a traceback and the record turns out to be missing three links.

The remainder of this article establishes that this is what has actually happened across four regimes, examines the one case where a supply base did move quickly and why, and sets out what a buyer should do differently. The short version of the recommendation is that the data supply should be scoped,

costed, and resourced before the platform is selected, and that the organization should establish early whether it possesses any leverage to compel the data it needs.

A further reason the mismatch persists is that the programme's internal reporting is accurate about the wrong thing. Milestones such as platform selected, integration complete, and internal events captured are all real achievements and all measurable, so a steering committee receives a progress report that is truthful and does not describe the state of the obligation. The only metric that describes the obligation is the proportion of upstream events actually arriving in usable form, and that metric is rarely defined at the outset because it requires knowing who the upstream participants are.

Organizations that add a single measure to their programme reporting, namely the percentage of required critical tracking events received electronically and in standard form from all participants in scope, obtain a different picture immediately. That figure typically starts far lower than the programme's completion percentage and moves much more slowly, and the divergence between the two lines is the most useful chart a traceability steering committee can look at. It shows the difference between what has been built and what has been achieved.

03 Three regimes, three multi-year slips

If the diagnosis is correct, it should be visible in the regulatory record, and it is. Three separate item-level regimes, in different sectors and different jurisdictions, set statutory compliance dates and each was extended by years. Figure 2 sets them side by side.

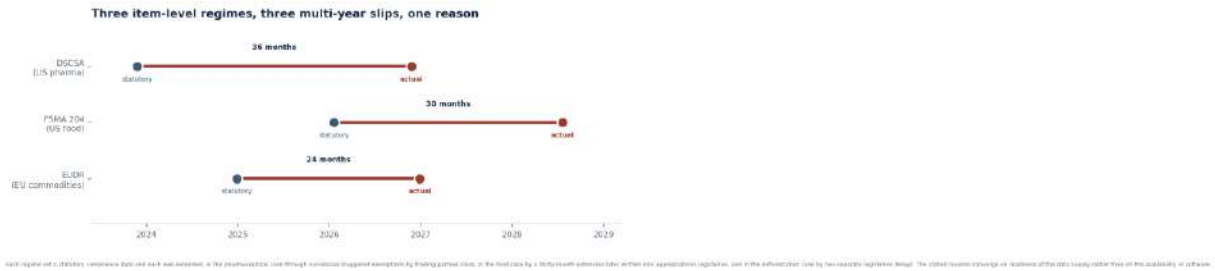

Figure 2. Three item-level regimes, three multi-year slips. Each set a statutory compliance date and each was extended, with stated reasons converging on readiness of the data supply rather than availability of software.

The pharmaceutical case is the oldest and most instructive. The enhanced drug distribution security requirements carried a statutory deadline of 27 November 2023, at which point trading partners were to exchange serialised transaction information electronically and be able to trace product at package level. In August 2023 the regulator announced a one-year stabilization period through November 2024, during which it did not intend to take action against parties continuing to work toward capability. That period was then followed by staggered exemptions extending to November 2026 for the smallest dispensers, which are examined in the next section.

The food traceability rule followed a similar path on a compressed timeline. The final rule became effective in January 2023 with a compliance date of 20 January 2026, requiring covered entities to maintain key data elements at defined critical tracking events and to produce a sortable electronic record within twenty-four hours of a request. In August 2025 the regulator published a thirty-month extension moving the date to 20 July 2028, and Congress subsequently directed in appropriations legislation that enforcement not occur before that date. A thirty-month extension to a rule finalised three years earlier is not a scheduling adjustment; it is a determination that the covered population could not comply.

The deforestation regulation shows the same pattern in a different jurisdiction. It entered into force in mid-2023 with application from 30 December 2024, requiring operators placing covered commodities on the market to conduct due diligence establishing that the goods were not associated with deforestation, supported by geolocation data for the plots of origin. The date was moved to December 2025 and then again to December 2026 for larger operators and June 2027 for micro and small enterprises, with a country benchmarking system published in May 2025. The stated reason for the second delay concerned the capacity of the information system to process the projected volume of due-diligence statements, which is a data-volume problem rather than a policy reversal.

It is worth noting what these three regimes have in common beyond the slippage, because the shared features explain why the pattern recurs. Each requires data at the level of an individual unit or lot rather than at the level of a shipment or a supplier relationship. Each requires that identity to survive multiple

changes of custody and, in the food and commodity cases, changes of physical form. And each places the compliance obligation on a party that sits downstream of most of the events it must record. Any future regime with those three properties should be expected to encounter the same constraint.

That expectation has planning value. An organization facing a newly announced item-level obligation can predict with reasonable confidence that the compliance date will move, that the movement will be justified on readiness grounds, and that the relief will favour smaller participants. This is not a reason to defer preparation, since the organization's own capability takes time to build regardless. It is a reason to sequence the work so that the supply-base development, which is the long pole, begins immediately, while the platform decision can reasonably wait until requirements have stabilised.

04 What the staggered exemptions reveal

The pharmaceutical extensions are worth examining in detail rather than as a single date change, because the structure of the relief is the clearest available evidence for this article's thesis. Figure 3 shows it.

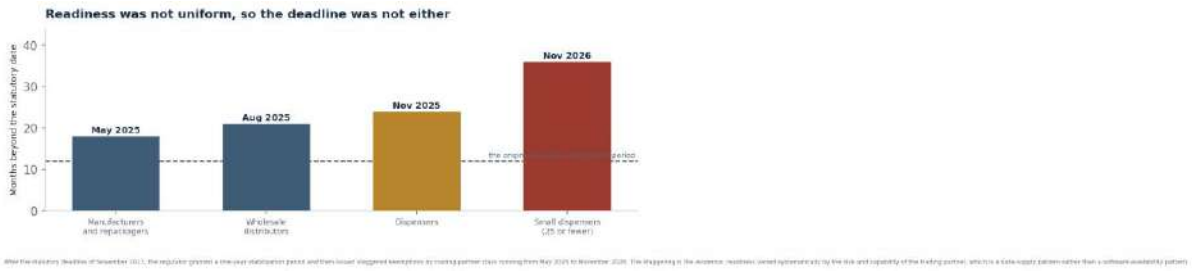

Figure 3. After the statutory deadline, the regulator granted a one-year stabilization period and then staggered exemptions by trading-partner class running from May 2025 to November 2026, in descending order of size and capability.

The relief was not granted uniformly. Manufacturers and repackagers, the largest and best-resourced participants, received exemption to May 2025. Wholesale distributors received August 2025. Dispensers received November 2025. Dispensers with twenty-five or fewer full-time pharmacists or pharmacy technicians, the smallest participants in the chain, received November 2026, a full three years beyond the statutory date. The ordering is precise and it maps exactly onto organizational capability rather than onto any difference in the underlying legal obligation.

This structure tells a buyer something no press release does. A regulator granting uniform relief is saying the whole system needs more time, which could reflect any number of causes including technology immaturity. A regulator granting relief in descending order of participant size is saying something much more specific: that the ability to comply correlates with the resources of the participant, that the large firms could do it and the small ones could not, and that the binding constraint is therefore capability distributed across the trading-partner population. That is precisely the data-supply constraint this article describes, expressed in the form of an exemption schedule.

The implication for a large organization is uncomfortable and important. Being ready oneself is not sufficient and never was. A serialised chain of custody requires every participant in the chain to exchange compatible data, so a manufacturer that achieved capability by 2024 still could not operate a fully electronic chain if its distributors and dispensers could not receive and forward the information. The largest participants effectively completed their own programmes and then waited three years for the rest of the chain, during which their investment produced no compliance benefit because compliance is a property of the chain rather than of any node in it.

A second implication concerns how organizations should read regulatory extensions generally. An extension is frequently treated as a reprieve, and the programme slows accordingly. Read correctly it is a diagnostic: it tells the organization where readiness actually sits, and the staggering tells it which parts of its own trading-partner population are the problem. An organization that responds to an extension by mapping its suppliers against the exemption classes learns exactly which of them the regulator has just

identified as incapable, and that is a considerably more useful response than deferring the programme by the length of the extension.

A third implication concerns budgeting and its timing. Organizations that treated the original statutory dates as fixed built business cases with benefits commencing at those dates and costs concluding shortly before them. When the dates moved, the cost was incurred on the original schedule while the compliance benefit was deferred by years, which produced exactly the pattern of stalled programmes and reduced executive confidence that makes the next phase harder to fund. The financial shape of these programmes deserves attention independent of their technical content.

A more robust approach separates the two components of the spend. Investment in the organization's own capability is worth making on the original schedule regardless of extensions, because it takes time and because it is a precondition for everything else. Investment in supplier enablement, which is the larger and slower component, should be phased against realistic readiness rather than against the statutory date, because spending it early against partners who cannot absorb it produces no capability and consumes the budget that will be needed when they can.

05 The regulator states the problem

The strongest single piece of evidence for this article's thesis is not an inference from the pattern of delays. It is the regulator's own explanation for the thirty-month food traceability extension, which states the data-supply constraint directly.

In explaining the extension, the agency observed that even entities that would themselves be ready by the original date faced challenges, in part because of their reliance on receiving accurate data from their supply chain partners, who might not be similarly situated. That sentence contains the entire argument. The constraint is not that the covered entities lack systems. It is that a covered entity's compliance depends on data originating with parties whose readiness it does not control and whose capability differs from its own. An entity can complete its programme perfectly and still be unable to produce a complete traceback, because the record it must assemble includes events that occurred elsewhere.

The structure of the food rule makes this dependency unusually explicit. It requires covered entities to maintain key data elements associated with critical tracking events, which include harvesting, cooling, initial packing, first land-based receiving, shipping, receiving, and transformation. For a leafy green reaching a retailer, most of those events occurred at a farm, a cooling facility, and a packer, none of which the retailer owns and several of which may have no relationship with the retailer at all. The retailer's obligation to produce a sortable record within twenty-four hours is therefore an obligation to have received data from parties several steps removed, in a compatible form, associated with a traceability lot code that survived every handoff.

The twenty-four hour requirement deserves specific attention because it converts a data-completeness problem into an operational one. An organization that holds most of the required records and can assemble the rest by telephoning suppliers over three days does not meet the standard. The record must be assembled and sortable within a day, which means it must already exist in electronic form before the request arrives. That requirement is what makes the upstream capability question binding rather than merely inconvenient, and it is why a programme that treats supplier data as something to be chased on demand will fail its first real test.

There is a further structural feature of the food rule worth drawing out, which is the traceability lot code and the point at which it is assigned. The rule requires that a lot code be established at defined points and carried forward, so that a unit received at the end of the chain can be associated with the specific harvest and packing events that produced it. Assignment is straightforward; persistence is not. Every subsequent handler must record the code it received, associate it with what it shipped, and transmit both, and any handler that fails to do so severs the association permanently.

The severing is what makes partial compliance so much less valuable than it appears. A chain in which nine participants of ten record and transmit correctly does not produce ninety percent of a traceback; it produces a traceback that stops at the tenth. This is why coverage statistics expressed as a percentage of suppliers can be actively misleading as a measure of programme health, and why the only meaningful test is whether a complete path can be reconstructed for a randomly selected unit.

06 Where the chain of custody actually runs

Understanding why supplier capability is the constraint requires looking at where in the chain the required events occur and what kind of organization is present at each point. Figure 4 sets out the structure.

Figure 4. The chain of custody runs through parties the buyer does not contract with. Capability declines with distance from the buyer, and at the origin the record is frequently on paper.

At the top of the chain sits the buyer's own operation, which is instrumented, standardised, and under direct control. Below it are direct suppliers, with whom the buyer has a contract and therefore leverage, and which in most sectors are large enough to have systems. Below that the position changes materially. The buyer has no contractual relationship with its suppliers' suppliers, capability varies from sophisticated to non-existent, and the buyer frequently does not know who they are. At the origin, in agricultural chains particularly, the participants may be growers, harvesters, and small processors operating on paper records or on no records at all.

The arithmetic of this is unforgiving. A chain of custody is complete only if every link captured and transmitted its events, so the completeness of the record is determined by the least capable participant rather than by the average. An organization with excellent systems, capable direct suppliers, and one paper-based participant at the origin has an incomplete record, and the incompleteness sits precisely at the point where the traceback most needs to reach. This is the opposite of how most quality and capability programmes work, where improvement at any point contributes to the whole, and it explains why traceability programmes can consume large budgets and still fail their first real test.

Two further features of the upstream population compound the difficulty. The first is economic: the participants least able to comply are also the smallest and least able to absorb the cost, and nobody has offered to fund them. A small grower asked to record and transmit event data at each handoff faces a real cost with no revenue attached, and the commercial relationship gives the distant buyer no mechanism to compensate for it. The second is structural: in many agricultural and extractive chains the intermediate aggregation steps deliberately mix product from many origins, which is efficient for trading and destroys the lot-level identity that traceability requires. Reconstructing identity through an aggregation point is not a data problem but a process change, and it costs money that the aggregator has no reason to spend.

This is why the practical recommendation later in this article begins with counting rather than with selecting. An organization that has classified its supply base by traceability capability, from fully instrumented through to paper-based, knows the size and shape of the problem it is actually solving. An

organization that has selected a platform knows what its own node will look like. The first exercise takes weeks and determines whether the programme can succeed; the second takes months and determines what the reporting will look like.

A practical consequence of this structure is that the buyer frequently cannot even enumerate the parties whose data it requires. Beyond the first tier the identity of participants is commercially sensitive, changes with the season or the market, and is in some chains deliberately obscured by intermediaries whose value proposition rests on holding those relationships. An organization that sets out to census its supply base to the origin will discover that a meaningful portion of it is unknown, and that discovery is itself a finding worth recording rather than an obstacle to the exercise.

Where identity cannot be established, the organization has two honest options. It can require its direct suppliers to disclose their sources as a contractual condition, which is achievable where leverage exists and which many buyers have not attempted. Or it can accept that the chain is opaque beyond a certain point and treat that opacity as a named risk with an owner, rather than as an implementation detail to be resolved later. What it should not do is build a programme plan that assumes the enumeration will happen without anyone having established that it can.

07 One retailer moved faster than three regulators

There is a case in the record that cuts against the pessimistic reading of everything above, and it is the most important evidence in this article because it identifies what actually works. Figure 5 sets it against the federal timeline.

Figure 5. Leverage moved the supply base. The deadline did not. A dominant grocery retailer required traceability event data with financial consequences from August 2025, while the federal rule covering similar data slipped to July 2028.

A major grocery retailer implemented its own supplier traceability requirement with effect from 1 August 2025. Suppliers were required to transmit advance ship notices carrying the relevant data elements, to apply serialised pallet labels, and to apply case-level labelling in a standard format. Non-compliance carried chargebacks, which is to say that a supplier failing to provide the data incurred an immediate financial penalty administered through the trading relationship rather than a distant regulatory risk. The supply base moved.

The comparison with the federal rule is stark and instructive. Both required broadly similar data elements at broadly similar points in the chain. The federal rule carried the authority of statute, the possibility of enforcement action, and a compliance date. The retailer requirement carried a chargeback. The federal date moved by thirty months; the retailer requirement took effect on schedule. The difference is not the strength of the obligation in principle but the immediacy and certainty of the consequence, combined with the fact that the party imposing it was economically indispensable to the suppliers concerned.

This case does two things to the argument of this article. It confirms the diagnosis, because it shows that supplier data capability is achievable and therefore that the failure of the regulatory programmes is not a technology limitation. And it complicates the pessimism, because it demonstrates that the constraint is not immovable: a party with sufficient leverage and a credible consequence can move a large supply base in a defined period. The variable is not capability in the abstract but whether anyone with leverage has required it in a way suppliers must respond to.

For an organization planning a traceability programme the practical question follows directly: what leverage do we have, and if we have none, whose can we borrow. A firm that is a significant customer of its direct suppliers has real leverage and should use it, through contract terms with defined data requirements and defined consequences. A firm that is a small customer has little, and its realistic options are to work through industry initiatives, to align its requirements with those of a dominant customer already imposing them, or to accept that its coverage will be partial and to focus its effort where the risk is concentrated. What does not work is issuing a requirement with no consequence attached and expecting a supply base to absorb a cost for it.

The retailer case also illustrates a subtler mechanism that regulators lack, which is standardisation by fiat. The retailer specified not merely that data be provided but exactly how: which label formats, which identifiers, which transmission method, at which points. Suppliers therefore had a single unambiguous target rather than a general obligation to be interpreted. Regulatory texts, by necessity, describe required outcomes rather than mandating a specific commercial format, which leaves each participant to determine its own implementation and multiplies the interoperability work across the chain.

An organization with leverage should learn from this and specify tightly. A requirement that suppliers provide traceability data will produce a heterogeneous collection of spreadsheets, portals, and file formats that costs more to consume than to have collected manually. A requirement that specifies the identifier scheme, the event structure, the transmission mechanism, and the timing produces data that arrives usable. Specificity is the difference between a supply base that has complied and a supply base whose compliance is worth having.

08 The enforcement regime that does bite

A fourth regime differs from the three examined so far in a way that is worth studying, because it did not slip and it has produced measurable consequences. Figure 6 shows the enforcement record.

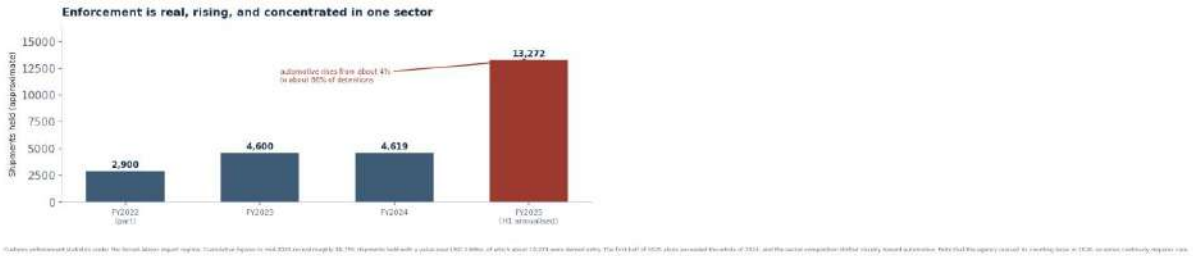

Figure 6. Customs enforcement under the forced-labour import regime. Cumulative figures to mid-2025 record roughly 16,755 shipments held with a value near USD 3.69 billion, of which about 10,274 were denied entry.

The forced-labour import regime took effect in June 2022 and operates through a rebuttable presumption: goods with specified regional links are presumed to be made with forced labour and are excluded unless the importer can demonstrate otherwise. Cumulative figures to mid-2025 record approximately sixteen thousand seven hundred and fifty-five shipments held with an aggregate value near three and seven tenths billion dollars, of which roughly ten thousand two hundred and seventy-four were denied entry and around five thousand seven hundred and eighty-three released. The first half of 2025 alone recorded around six thousand six hundred detentions, exceeding the whole of the preceding year.

Two features of the composition are instructive. The sector concentration shifted dramatically, with automotive rising from a small fraction of detentions to the substantial majority within a year, which indicates that enforcement attention moves and that a sector currently untouched should not assume it will remain so. And the origin pattern shows detentions and denials associated with third countries running well ahead of those associated with the primary jurisdiction, which reflects the fact that goods reach the border through intermediate manufacturing and that establishing origin through those steps is exactly the traceability problem this article describes.

The reason this regime bites where the others slipped is structural. Compliance is enforced at the border by an agency that can simply hold the goods, and the consequence of an incomplete record is immediate, commercial, and severe: the shipment does not enter. There is no extension mechanism from the importer's point of view, because the presumption operates on the specific consignment in front of the officer. This is the regulatory equivalent of the retailer chargeback examined in the previous section, and it produces the same behaviour, which is that importers invest in tracing their upstream chains because the alternative is losing the goods.

A methodological caution belongs with these figures. The agency revised its counting basis in 2026 toward a per-transaction measure, which produced a substantially higher cumulative shipment count against a similar value figure. Series continuity therefore requires care, and comparisons across the revision should be made on value rather than on count. This is a common hazard with enforcement statistics and it is worth stating rather than reporting the higher figure as growth.

The composition shift toward automotive carries a lesson that generalises beyond this regime. A sector that had accounted for a small fraction of detentions became the substantial majority within roughly a year, which means that firms in that sector experienced a step change in enforcement exposure without any change in their own sourcing behaviour. Enforcement attention is allocated by an agency responding to intelligence, political direction, and the results of its own prior work, and it moves faster than a supply base can be restructured.

For a firm in a sector not currently under scrutiny, the practical implication is that the time to establish upstream visibility is before attention arrives rather than after. Firms that were already able to document their sub-tier sourcing when the automotive shift occurred were able to respond to detentions with evidence; firms that were not faced the choice between holding goods at the border while they investigated and abandoning the consignment. The capability took months to build and was needed within days.

09 The evidence base, and who publishes it

An organization building a business case for a traceability programme will encounter quantified benefit claims, and the provenance of those claims requires attention before any of them enters a model. Figure 7 sorts the figures in circulation.

Figure 7. Five figures in circulation and what each actually is. Almost every quantified benefit claim originates with a party that sells traceability or compliance software; only the regulator-authored study is a measurement.

The claims fall into a consistent pattern. A compliance vendor reports that traceability narrowed a recall from six days of production to one, which is an anecdote from a single implementation with no counterfactual. An enterprise software vendor reports that root-cause investigation time drops from days to minutes and that traceback effort falls by eighty percent, which are marketing claims with no disclosed method. An average recall cost figure of around ten million dollars circulates widely, attributed to a trade association and repeated by vendors, with the underlying method not visible. Each of these may be broadly right and none is evidence in the sense a business case requires.

Against these sits one category of source that is different in kind. Regulator-authored, peer-reviewed analysis of three consecutive leafy-green outbreak investigations in 2018 and 2019 documented four hundred and seventy-four illnesses, two hundred and fifteen hospitalisations, and five deaths, and examined why the traceback exercises could not identify sources quickly enough to limit exposure. That is a measurement, published by the agency that conducted the investigations, in a peer-reviewed venue, describing what actually happened. It is the strongest available evidence for why item-level traceability matters and it is a public-health argument rather than a cost-saving one.

The standards position deserves a note of its own. The data formats and identifier schemes underpinning most of these regimes come from a standards body funded by its members, which include the largest participants in the chains being regulated. This is the normal way industry standards are developed and it is not a criticism. It does mean that published adoption figures and readiness assessments originating with the standards body describe a population with an interest in the standard's success, and they should be read as such. Where an organization needs to know how many of its own suppliers can actually transmit compatible data, the reliable method is to ask them and test it rather than to consult a published adoption rate.

The practical guidance follows directly. A traceability business case should be built on the regulatory obligation and its consequences, which are certain, and on the organization's own risk exposure, which it can assess. It should not be built on vendor-published efficiency claims, because those cannot be verified and because the programme will be judged on whether it produces a complete traceback rather

than on whether it saved investigation time. Where a benefit claim is used, its source and interest should be stated in the paper, which takes one clause and prevents a marketing figure from acquiring the standing of a finding.

One further category of claim deserves separate treatment because it is the most consequential and the least examined, which is the assertion that a traceability programme will reduce the scope of a recall. The logic is sound in principle: better records permit a narrower withdrawal. Whether it holds in a specific case depends entirely on whether the records reach the granularity of the contamination, and a chain that traces to the packer but not to the field will support a narrower recall than no traceability at all and a considerably wider one than the vendor example implies.

An organization evaluating this benefit should therefore ask what level of granularity its own chain will actually achieve rather than what the technology permits. The answer determines the size of the recall the firm could execute, which is the quantity that matters commercially. Stating that answer explicitly in the business case, with its dependence on upstream participation made visible, produces a more defensible number and has the useful side effect of directing attention to the upstream gap that determines it.

10 What the standards do and do not settle

A recurring assumption in these programmes is that adopting the relevant data standards resolves the interoperability problem, and it resolves an important part of it while leaving the harder part untouched.

What the standards do settle is materially valuable. They define identifier schemes so that a product, a location, and a logistics unit have unambiguous references. They define event structures so that a shipping event captured by one party can be interpreted by another. They define exchange formats so that systems from different vendors can transmit records without bespoke integration. Without this layer, a chain of custody spanning many organizations would require pairwise agreements between every trading partner, which does not scale. The standards work is a precondition for any of this to function.

What the standards do not settle is whether the event was captured at all, whether it was captured accurately, and whether the party captured it for the right unit. A standard specifies how to express that a particular lot was received at a particular location at a particular time; it does not cause anyone to record that fact, and it cannot detect that the fact recorded was wrong. In chains where the upstream participants are small, where recording is manual, and where the person doing the recording has other priorities, the accuracy of the captured events is the dominant source of error and no format specification addresses it.

A related limitation concerns the traditional one-up, one-down structure of much traceability practice, in which each party records who it received from and who it shipped to. This is straightforward to implement and it produces a complete chain only if every link participates, because the chain is reconstructed by walking the links sequentially. A single non-participating or inaccurate link breaks the walk, and the party attempting the traceback discovers this only when it tries. Regimes that require lot-level identity to persist through transformation are attempting to address this, and they impose the requirement on exactly the aggregation and processing steps that have the strongest economic reasons to mix product.

There is a governance question underneath the standards discussion that organizations rarely ask, which is who inside the firm owns the accuracy of the captured events. Ownership of the platform sits with technology, ownership of the obligation sits with compliance or regulatory affairs, and ownership of the supplier relationship sits with procurement. Accuracy of the event data sits with none of them, and consequently with nobody. That gap is where the failures accumulate, because an inaccurate event is not detected by any of the three functions until a traceback fails.

The remedy is to name an owner for data accuracy explicitly, with the authority to test and the mandate to escalate. This is not an elaborate structure; in most organizations it is one accountable individual with a quarterly test and a route to the executive who can change a supplier relationship. What it replaces is a situation in which three functions each believe another is verifying the data, which is the ordinary state of affairs in programmes of this kind.

11 The fairness case: the mandates are working

This article has catalogued delays and structural obstacles, and a reader who concluded that item-level traceability mandates are futile would be drawing a conclusion that the evidence contradicts and that this article does not intend.

The strongest point is the public-health case, and Figure 8 states it. Three consecutive leafy-green outbreak investigations produced four hundred and seventy-four illnesses, two hundred and fifteen hospitalisations, and five deaths, and the regulator's own peer-reviewed analysis identified the traceback exercise as the bottleneck: investigators could not establish the source quickly enough to limit exposure. Set against roughly forty-eight million foodborne illness cases and around three thousand deaths annually in the United States, the case for being able to identify a contaminated lot within a day rather than a fortnight is not a compliance argument at all. It is the reason the rule exists, and no analysis of implementation difficulty diminishes it.

Figure 8. Regulator-authored, peer-reviewed analysis of three leafy-green outbreak investigations. Investigators could not identify sources quickly enough to limit exposure, and the traceback exercise itself was the bottleneck.

The second point is that mandates demonstrably move industries that nothing else moves. The retailer case in Figure 5 shows a supply base responding within months to a requirement with a consequence attached, and the forced-labour regime shows importers investing in upstream tracing because goods are being held at the border. Both are mandates. The lesson is not that mandates fail but that mandates without immediate consequence fail, which is a design point rather than an objection to the instrument. A regulator with an enforcement mechanism that bites at the point of transaction gets compliance; one relying on a distant compliance date gets extensions.

A third point in fairness concerns the extensions themselves, which this article has used as evidence of failure and which can be read more charitably. A regulator that extends a deadline rather than enforcing against a population that cannot comply is making a reasonable judgment about proportionality, particularly where the unprepared population consists of small entities that would be driven out by enforcement. The staggered structure examined in Figure 3 is an attempt to hold the capable to the schedule while giving the incapable time, which is a more sophisticated response than either blanket enforcement or blanket delay. It is slower than the statute intended and it may well be the correct policy.

A fourth point deserves emphasis because it qualifies the article's own framing. Where these regimes have been implemented, they have delivered. Pharmaceutical serialisation has made counterfeit introduction materially harder in regulated distribution channels. Import enforcement has changed

sourcing behaviour in affected sectors, visible in the shifting composition of detentions. Deforestation due diligence has forced operators to obtain geolocation data that did not previously exist. The programmes are late and they are producing the capability they were designed to produce. The criticism in this article is directed at how organizations scope and resource these programmes, not at whether the obligations are worth meeting.

A fifth point in fairness concerns what these regimes have made visible that was previously invisible. Before the deforestation obligation, most operators placing covered commodities on the market could not have said with confidence which plots of land their material came from, and many discovered on investigation that their supply chains reached places they had not known about. That discovery is itself valuable independent of the compliance outcome, because a firm that does not know where its material originates cannot manage any risk associated with the origin, whether regulatory, reputational, or operational.

The same applies to the forced-labour regime, where the requirement to rebut a presumption forced importers to construct sub-tier maps that they had never previously needed. Several of the firms that did this work reported finding concentrations and dependencies they had not known they held. A mandate that produces that kind of visibility has delivered value even where the compliance burden was heavy and the timeline uncomfortable, and any assessment of these regimes that counts only the cost of compliance has counted only one side.

12 Buy the constraint, not the platform

The constructive principle follows from the diagnosis: scope, cost, and resource the data supply before selecting a platform, because the data supply is the binding constraint and the platform is not.

In practice this begins with counting. The organization should classify its supply base by traceability capability, and the classification does not need to be elaborate: participants with instrumented systems that can already transmit standard event data; participants with systems that would require configuration or integration; participants with basic electronic records but no event capability; and participants operating on paper. That census can be completed in weeks through the existing supplier relationships, and it immediately reveals the size and shape of the problem. Most organizations that run it find the distribution considerably worse than assumed, particularly beyond the first tier.

The second element is to establish where the leverage is. For each segment of the supply base, the organization should ask whether it is a large enough customer to impose a requirement with a consequence, whether an industry initiative or a dominant customer is already imposing compatible requirements it can align to, or whether it has no practical means of compelling the data. This produces a realistic map of achievable coverage, which is a far more useful planning artefact than a target of complete coverage that the organization has no mechanism to reach. Where leverage is absent, the honest conclusion is that coverage will be partial, and the effort should be concentrated where the regulatory or safety exposure is greatest.

The third element is to fund the constraint. Where the participants that cannot comply are small and the cost of compliance is real, the organization has a choice between requiring the capability without funding it, which produces slow and unreliable adoption, and contributing to it, which is unusual and considerably more effective. Contributions can take several forms: providing the labelling equipment, providing a simple capture application, absorbing the integration cost, or paying a price premium tied to compliant data. Organizations that treat upstream capability as something to be procured rather than demanded generally achieve coverage faster, and the cost is small relative to a platform programme.

A fourth element concerns sequencing within the supply base rather than across it. Not all upstream participants matter equally: the commodities subject to the strictest obligations, the suppliers handling the highest volumes, and the points in the chain where identity is most likely to be lost are where effort produces the most coverage per unit of work. An organization that segments on this basis can produce meaningful traceback capability for its highest-exposure products long before universal coverage is achievable, which is both better risk management and better use of a finite budget.

This segmentation also produces a defensible position in the event of regulatory attention. An organization that can demonstrate a deliberate, risk-based sequencing plan, with the highest-exposure categories complete and a documented schedule for the remainder, is in a materially stronger position than one that has spread its effort uniformly and achieved partial coverage everywhere. Regulators generally respond better to evidence of prioritised progress than to uniform incompleteness, and the former is also more likely to prevent the incident that matters.

13 A traceability protocol, and a scoring rubric

The principles above combine into a protocol and a rubric a governance body can apply to judge whether a traceability programme is addressing its actual constraint. Figure 9 sets out the discipline.

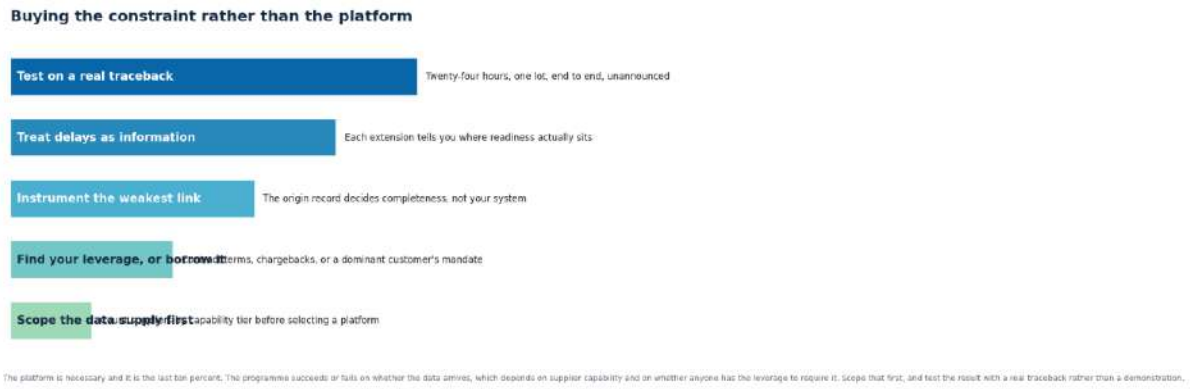

Figure 9. Buying the constraint rather than the platform: scope the data supply first, find or borrow the leverage, instrument the weakest link, read delays as information, and test with a real traceback.

The protocol runs as follows. Census the supply base by traceability capability before selecting any platform, extending the census beyond the first tier for the commodities where the obligation actually bites. Establish where leverage exists and where it does not, and set achievable coverage targets accordingly rather than declaring universal ones. Where participants cannot comply and cannot fund compliance, decide explicitly whether to fund it or to accept partial coverage in that segment. Treat each regulatory extension as a diagnostic about where readiness sits rather than as a reprieve. And test the resulting capability with an unannounced end-to-end traceback against the actual regulatory standard rather than with a system demonstration.

A scoring rubric

The dimensions below distinguish a programme addressing its constraint from one implementing a platform.

Dimension	Addressing the constraint	Implementing a platform
First activity	Supply base capability census	Vendor selection
Scope of view	Beyond tier one to the origin event	Own systems and direct suppliers
Leverage	Mapped, with consequences defined	Assumed, or not considered
Upstream cost	Funded or explicitly accepted as a gap	Assumed to be the supplier's problem
Response to delay	Read as a readiness diagnostic	Programme slows by the length of the extension
Evidence base	Regulatory text and own testing	Vendor efficiency claims
Verification	Unannounced end-to-end traceback	System demonstration

A programme scoring in the left column knows how complete its chain of custody actually is and where it breaks. A programme scoring in the right column has an implemented platform and will discover the state of its chain when a regulator or a recall requires it. The rubric does not reduce the work; the census and the leverage mapping are additional effort at the front of the programme. It relocates the work to the part of the problem that determines whether the obligation is met.

Two additions to the rubric are worth making where the organization operates across multiple regimes simultaneously, which is increasingly the normal case. The first is whether the programme has identified the overlap between obligations, since the underlying data elements required for food traceability, forced-labour rebuttal, and deforestation due diligence share a common core of origin, custody, and volume records. Organizations that treat each obligation as a separate programme build the same supplier capability three times and pay for it three times.

The second is whether the supplier-facing requirement is unified. A supply base receiving three separate data requests from three functions of the same customer, in three formats and on three schedules, will comply slowly and inconsistently with all of them. A single consolidated requirement, specified once and covering the union of what the obligations demand, is materially easier for a supplier to implement and therefore materially more likely to be implemented. This is an internal coordination problem that shows up as an external capability problem, and it is within the buyer's control.

14 Test it with a real traceback

Of everything in the protocol, one verification step is worth separate treatment because it is the only one that establishes whether the programme actually works, and because almost nobody performs it.

The test is straightforward to specify. Select a unit or lot at random from finished inventory, without notice to the participants involved, and attempt to assemble the complete chain-of-custody record to the standard the regulation requires, within the time the regulation allows. For a food traceability obligation that means a sortable electronic record within twenty-four hours covering every critical tracking event back to the origin. For a pharmaceutical obligation it means the transaction information and statement at package level through every trading partner. The test is passed only if the record is complete, accurate, and produced within the time limit.

What such tests reveal is consistent and useful. The organization's own records are generally excellent. The first tier is generally adequate. The record thins or breaks somewhere upstream, frequently at an aggregation or transformation point where lot identity was not preserved, and the missing segment can only be reconstructed by telephone over several days, which fails the standard. Organizations that run this test learn precisely which links are weak, which is exactly the information the census in the previous section is designed to produce and which the test verifies independently.

The unannounced element matters for the same reason it matters in any audit. A traceback exercise conducted with notice assesses a chain whose participants have prepared, and preparation for a traceback consists largely of assembling records that should already have been in transmissible form. The test is specifically designed to establish whether the record exists before it is requested, which is the regulatory standard, so giving notice defeats the purpose entirely. This is uncomfortable to arrange with trading partners and it is the only version of the test that produces information.

The test should also be repeated rather than performed once at go-live. Supply bases change, suppliers substitute sub-suppliers, seasonal sourcing shifts origins, and a chain that was complete in one quarter may not be in the next. An organization that runs a traceback test quarterly on a randomly selected lot has a continuous measure of a capability that otherwise degrades invisibly, and it will discover breaks in the ordinary course rather than during a recall. The cost is a few days of work per quarter, which is trivial against the exposure it measures.

One refinement to the test is worth adopting where the regulatory obligation permits it, which is to run it against the specific commodity or product where the exposure is greatest rather than against a uniformly random selection. Random selection produces an unbiased estimate of overall capability, which is useful for reporting. Selection weighted toward the highest-risk categories produces information about the chains where a failure would matter most, which is more useful operationally. Most organizations should do both, and should not confuse the two.

The results should also be recorded and retained rather than treated as an internal exercise. A documented history of periodic traceback tests, with the failures and the remediation, is evidence of a functioning compliance programme in a way that a platform implementation certificate is not. Where a

regulator or a customer asks how the organization knows its chain of custody works, a test log answers the question directly and a system architecture diagram does not.

15 Conclusion: leverage, not licences

Item-level traceability is among the clearest examples in enterprise technology of a problem that is procured in one place and located in another. The obligation concerns an entire chain of custody. The purchase is made by one organization at one point in that chain. The platform that organization buys will work, its own records will be excellent, and its programme will report complete. Whether the obligation is actually met depends on parties it does not own, mostly does not contract with, and in agricultural and extractive chains frequently cannot identify.

The regulatory record is unusually explicit about this. Three separate item-level regimes in two jurisdictions set statutory dates and each slipped by years. The pharmaceutical relief was granted in descending order of participant size, which is a regulator stating that capability rather than intent is the constraint. The food traceability extension of thirty months was justified on the ground that entities ready themselves still depend on receiving accurate data from partners who might not be similarly situated, which is the thesis of this article in the agency's own words. The deforestation delay turned on the volume of due-diligence statements the system would have to process. None of these was a technology failure.

What works is visible in the same record. A grocery retailer required traceability event data from its suppliers with chargebacks attached and the supply base moved within months, three years ahead of the federal rule covering similar ground. A border enforcement regime that holds goods produced sustained investment in upstream tracing across affected sectors. Both are mandates with an immediate, certain, commercial consequence delivered by a party the supplier cannot avoid. The variable that predicts whether a supply base becomes traceable is not the strength of the obligation in law but whether someone with leverage has attached a consequence to it.

For an organization planning or reviewing such a programme the implications are practical. Census the supply base by capability before selecting a platform, and extend the census past the first tier to the point where the events actually occur. Map where leverage exists, set coverage targets to what leverage can achieve, and be explicit about the segments where it cannot. Where the participants who cannot comply also cannot fund compliance, decide whether to fund it rather than issuing a requirement that will not be met. Read each extension as a diagnostic rather than a reprieve, and use the staggering to identify which of your own partners the regulator has just classified as incapable. And test the result with an unannounced end-to-end traceback against the actual standard, repeated quarterly, because that is the only exercise that distinguishes a capability from an implementation. The mandate will not be met by the system that was purchased. It will be met, or not, by the data that arrives.

There is a final observation worth making about how these programmes are governed, because it determines whether the analysis above can be acted upon. A traceability programme reported through a technology steering committee will be assessed against technology milestones, and will therefore continue to report progress while the constraint remains untouched. The same programme reported to an operating or risk committee, against the measure of how much of the required upstream data actually arrives, will surface the constraint within a quarter.

That single change in reporting line and metric is the least expensive intervention available and frequently the most effective. It does not require additional budget, new technology, or a supplier negotiation. It requires the organization to define the measure that describes the obligation rather than the effort, and to put that measure in front of the people who can allocate the leverage needed to move it.

16 Methodology, caveats, and sources

Methodology

- This article draws on primary regulatory sources, legislative and rulemaking records, agency enforcement statistics, and regulator-authored peer-reviewed research, current to mid-2026. Supply Chain Research is independent and accepts no payment from the traceability software vendors, compliance advisers, or standards bodies discussed.
- Benefit claims published by parties that sell traceability or compliance software are identified as interested sources and are not used as evidence. The standards body whose formats underpin these regimes is funded by its members, and its published adoption material is treated accordingly.

Caveats

- Regulatory dates and requirements described here were current at the time of writing and are subject to further amendment, extension, and legislative action. Organizations with compliance obligations should consult the current primary sources and their own advisers rather than relying on this summary.
- Enforcement statistics under the forced-labour import regime were revised to a per-transaction counting basis in 2026, which produced a materially higher cumulative shipment count against a similar value figure. Comparisons across that revision should be made on value rather than count.
- Recall-time and investigation-time savings figures in circulation originate with compliance and enterprise software vendors, are anecdotal or unmethodologised, and are reported here only as an illustration of the evidence base rather than as measurements.
- The retailer mandate described is drawn from published supplier guidance and secondary reporting rather than from the retailer's own regulatory filings, and details of scope and enforcement may differ from the summary here.
- Figures 1, 4, and 9 are conceptual illustrations of structure rather than measured data. Figure 2 presents statutory and revised dates rather than a continuous series.
- This article addresses the data-supply constraint in traceability programmes. It does not assess the technical merits of any particular platform, standard, or vendor.

Sources

[1] US Food and Drug Administration. Compliance policies establishing the one-year stabilization period for electronic drug distribution systems.

[2] US Food and Drug Administration. Waivers and exemptions beyond the stabilization period.

[3] US Food and Drug Administration. Food traceability final rule: requirements for additional traceability records for certain foods.

[4] Federal Register. Requirements for additional traceability records for certain foods: compliance date extension.

[5] US Food and Drug Administration. Constituent update on the proposed extension of the food traceability compliance date.

[6] Council of the European Union. Press release on the further postponement of the deforestation regulation.

[7] US Customs and Border Protection. Forced labour prevention act enforcement statistics.

[8] US Food and Drug Administration authors, PubMed Central. Overview of traceback investigations and three case studies of outbreaks linked to leafy greens.

Additional context drawn from European Commission material on deforestation regulation country benchmarking and simplification, from published retailer supplier-traceability guidance, and from vendor and standards-body publications which are identified as interested sources wherever used. This article is analysis, not legal, regulatory, or compliance advice, and its conclusions should be validated against the current primary sources and your own advisers before any decision.