WHITE PAPER

The execution record

Visibility tools read the record. Execution tools write it. The first mile is won or lost in the writing, and this paper is about what a record has to hold before anything can report on it.

Every import team has a dashboard. Ask it what is in carton 0042 and it goes quiet. Ask whether the label on that carton matches the record the retailer received, and it points at a document someone else maintains. Ask who last touched the carton and when, and the trail ends at a login. The dashboard is not broken. It is reading a record that was never written.

This paper is about the difference between the layer that reads a supply chain and the layer that writes it, why the difference decides the first mile, and how to test any system against it. It is written for the people who carry the fine and the deadline, not for a category debate.

1. The question the dashboard cannot answer

Visibility products are good at the questions they were built for: where is the container, when will it arrive, what changed since yesterday. Those questions are about movement, and movement is observable from the outside. A carrier event, an AIS ping, a terminal feed and a vessel schedule are enough to answer them.

The first mile asks different questions. They are questions about composition and provenance: what exactly is inside the carton, which PO lines it serves, which destination it is bound for, which label was printed onto it, and which document told the retailer so. None of those facts can be observed from outside the goods. They have to be written down at the moment they are made: at packing, at booking, at consolidation, at handover.

A system that did not write them cannot report them. It can join other systems that did, and it can reconcile. Reconciliation is the tell. When a supply chain answer ends with "let me check with the supplier", the record lives in people and the software is a reader.

2. Read versus write

The distinction is not about features. It is about where the fact is created. A read-only layer receives facts: from a carrier, a customer, a forwarder, an ERP extract. An execution layer creates facts as a side effect of the work: packing a carton mints its SSCC, booking against POs writes the link, building a container writes the consolidation, and a leg writes its own events.

Both kinds of system can display the same screen. Only one can answer for it under audit. The read-only layer can show a carton trace that is 98 per cent complete, and the two per cent is exactly the carton in

dispute. The execution layer cannot have a carton without a record, because the record is what the carton is.

This is why data quality work on a read-only stack is endless. Every gap is patched by asking someone to send a file. The gap returns with the next supplier, the next customer portal and the next quarter, because the same fact keeps being created somewhere the system does not reach.

3. The object chain is the record

A first-mile record is not a flat table of shipments. It is a chain of linked objects, and the links are the point. Shipeezi models it in the order the goods move: a PO line, an SSCC carton that holds its SKU contents as rows, a supplier booking that carries a mode and the cartons it covers, a container built from bookings, and an ordered chain of legs, each with its own mode and operator.

Each object carries the fields it owns. Origin and plot geometry belong at the SKU line, because that is where the goods came from. Emissions belong at the leg, because emissions are a property of movement. Title belongs to the document that transfers it. A container owns its vessel, voyage and seal; a pending supplier booking does not. Getting the ownership right is what lets compliance evidence be emitted from the same record instead of reconstructed beside it.

The links are what make the five questions answerable. A carton knows its lines. A booking knows its cartons. A container knows its bookings. A leg knows its container. Remove any link and the record becomes a set of tables that happen to mention the same SKU.

4. Three gaps a read-only stack creates

The scattered record

In a typical import operation the same fact lives in five places: the ERP holds the PO, the supplier's spreadsheet holds the pack, the forwarder's system holds the booking, the CFS holds the receive, and a portal holds a copy of all four. Completeness is a property of the whole, and no one owns the whole. A question that crosses a boundary becomes a meeting.

The cost is not only administrative. Scattered records make every exception ambiguous. When a carton is short, each system can prove it did its part, and none can prove what actually happened between them. The investigation consumes the margin on the shipment.

The label gap

A scan proves that a label was seen. It does not prove that the label matches the record, or that the contents match the label. Reader systems inherit this gap. They can validate what they receive, but they cannot guarantee the identity they are validating, because the label was minted somewhere else.

When the system mints the SSCC at packing, the label and the record become one fact. The 856 can be generated from the same data and match the physical barcode by construction. The carton is not described by the record. It is the record, printed.

The compliance retrofit

Compliance regimes ask for fields at specific nodes. CBAM asks about emissions, which belong to a leg. UFLPA asks for origin and supply chain due diligence, which belongs at the SKU line. EUDR asks for plot geometry, which belongs even further down. A system that holds a shipment summary cannot answer any of them without a backfill, and a backfill against live import data is a data project with a regulator attached.

The practical test is simple: can you emit the evidence for a shipment from the record that ran it, on the day it is asked for, without contacting a supplier? If the answer is yes, compliance is a property of the record. If it is no, compliance is a project.

5. The five-question test

Any first-mile system can be evaluated with five questions. Ask them of the vendor and of yourself.

- What is inside carton X, at SKU grain, and which PO lines does it serve?
- Does the label on carton X match the record and the ASN that was sent, and can you show that it always did?
- Who touched carton X, in what role, and when, and can you show the trail without a support ticket?
- Can you produce the compliance evidence for shipment Y from the same record, on demand?
- When a chargeback arrives, can you point at the field that caused it, and fix the source rather than the document?

A reader passes the first question for containers and fails it for cartons. An execution record passes all five, because each answer is a projection of the same chain.

6. What changes when the record is real

Adoption changes first. When the record is shared, a supplier, forwarder or CFS operator is not being asked to maintain a second copy in a portal. They work the same POs and cartons, scoped to what they are entitled to see. The invitation is cheap and the data stays one version. Most portal fatigue is not resistance to software. It is resistance to retyping.

Chargebacks change next. Pre-ship validation has something to validate: a label that matches the record, a carton that matches the ASN, a timing rule that can be checked before departure. The fine is prevented rather than disputed. Finance stops funding a reconciliation that should not exist.

Agentic work changes last, and it is the reason the record matters now. An agent is only as good as the facts it is allowed to act on. An agent reading a scattered record optimises against stale numbers; an agent reading an execution record acts on the same fact every other party sees, with an audit trail behind it. Autonomy without a ledger is a liability. The ledger is the product.

7. How to evaluate a system

Ask eight questions, and insist on a demonstration rather than a description.

- Where is the carton record created, and by whom?
- Does the system mint the label it validates, or receive it?
- What owns the PO line, the carton, the booking, the container and the leg, and what are the links between them?
- Which fields live at which node, and why?
- Can the 856 be generated from the pack record as a function, and tested before transport?
- What does a supplier see, and are they retyping anything?
- What does the audit trail record, and can it be exported?
- Can the same record emit compliance evidence without a data project?

If the answers converge on one chain, with one owner per field and one version of the truth, the first mile has a record. If they converge on integrations, you are buying reconciliation with a nicer interface.

8. Closing: the first mile writes the facts

The first mile is not a visibility problem, because visibility is downstream of writing. It is not a collaboration problem either, because collaboration without a shared record is copying. It is a record problem: who creates the fact, at which node, with which identity, and what can be proven afterwards.

Fix the record and the dashboards get honest for free. Leave it scattered and every new tool, including every AI agent, inherits the gap and industrialises it. That is the choice the first mile is asking you to make.