White Paper

Labeling as a Controlled Process

Part 1: Introduction - Why labeling deserves the same discipline as any other controlled process, and what to do about it

Author: Ján Uhrinovský

Reviewer: Gerard Ipskamp

Date Published: Mar/25/2026

Introduction

Every product that leaves a factory carries a label or mark. At the point of shipment, that label is the primary carrier of product identity - it tells customers, distributors, regulatory authorities, and traceability systems what the product is, where it came from, what version it is, and what rules apply to it.

When the label is wrong, the consequences are not minor. At best, there is a reprint and a delay. At worst: a shipment hold, a customer complaint, a regulatory investigation, or a recall. And the investigation that follows - trying to prove which label version was on which unit on which day - can take weeks if the evidence was never properly designed.

Labeling is not owned by one function. It sits at the crossroads of four: Quality, Packaging Engineering, Operations, and IT/OT. When no function owns the whole picture, each one solves its own piece locally. Over time, those local solutions accumulate into something fragile: variation between sites, informal knowledge that walks out the door when people leave, and an evidence trail that can only be reconstructed, not retrieved.

The uncomfortable truth:

In many organizations, labeling is still managed as a peripheral activity - something printers do at the end of the line, with templates updated informally and changes applied whenever someone has time. For a process this critical to product identity, that is a surprisingly fragile approach.

This methodology is most directly relevant where labeling complexity is highest: multi-site manufacturing operations, regulated product categories, environments with customer-specific label requirements, high change velocity across products or markets, and any production context where traceability must withstand audit or investigation. If your operation has more than one site, more than one market, or a history of labeling incidents that were harder to investigate than they should have been, this paper describes the structural cause and a practical remedy.

The four ways labeling goes wrong

Across manufacturing environments, the same four failure patterns appear repeatedly. They are not caused by individual mistakes - they are structural, driven by how the work is organized. The four failure modes below reflect practitioner observations across multi-site, regulated, and high-change-velocity environments. Recognizing them is the first step to fixing them.

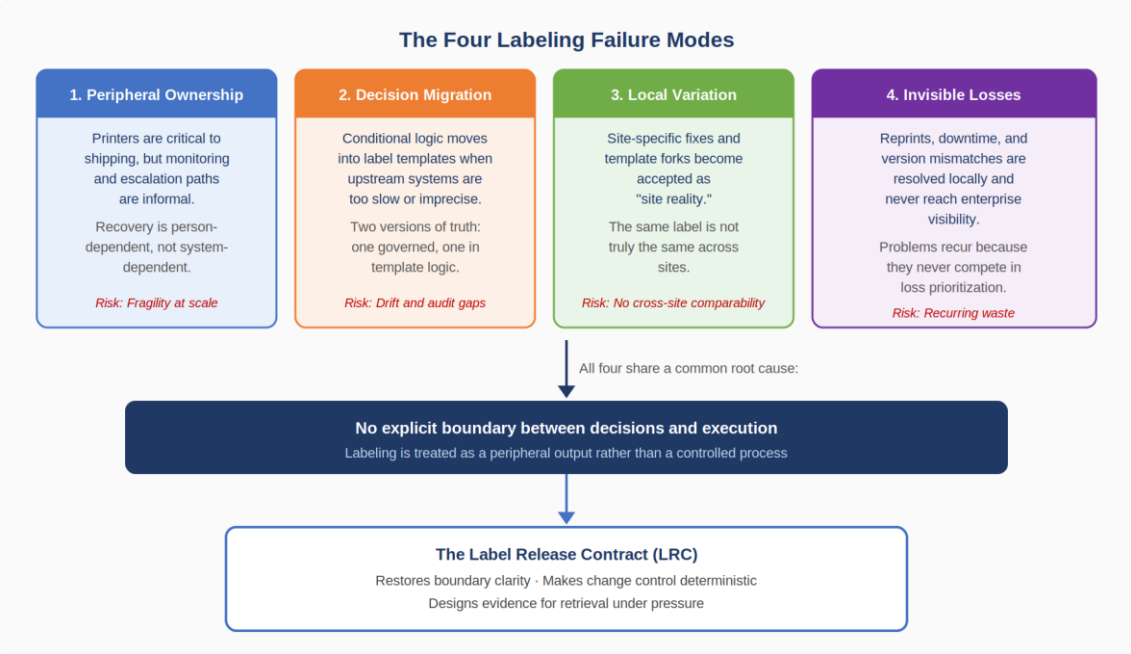

Figure 1: The Four Labeling Failure Modes

1. Peripheral ownership - recovery depends on people, not systems. Printers and label applicators are critical to shipping, but the monitoring, escalation paths, and recovery procedures around them are often informal. When something goes wrong, recovery depends on whoever is there and what has worked before. The plant keeps shipping - but only because experienced people carry complexity that the system does not.

2. Decision migration - logic moves into templates. When upstream systems do not provide the right data at the right time, teams embed conditional logic directly into the label template. This creates a hidden problem: the template becomes the real decision-maker, and no one can easily audit what rules it contains. Over time, the organization ends up with two versions of the truth - one in the governed systems, and one embedded in template logic.

3. Local variation - fixes become features. Once a site-specific workaround is accepted as standard practice, it stops being a problem to solve and becomes a fixed feature of the site. Improvements developed at one site do not translate to another. Cross-site comparisons become meaningless, and rollouts become very difficult.

4. Invisible losses - problems never compete for attention. Reprints often do not appear as waste in operational data. Labeling downtime gets buried under generic stop categories. Version mismatches are resolved locally and never reach enterprise visibility. Because these problems are not measured, they never compete for resources in loss reviews.

What these four patterns have in common:

All four stem from the same root cause: no explicit boundary between the decisions that determine what goes on a label and the execution process that produces it. The next section describes the governing mechanism that draws that boundary.

The Label Release Contract

The methodology draws a hard line between two things that most labeling environments blur together: decisions and execution.

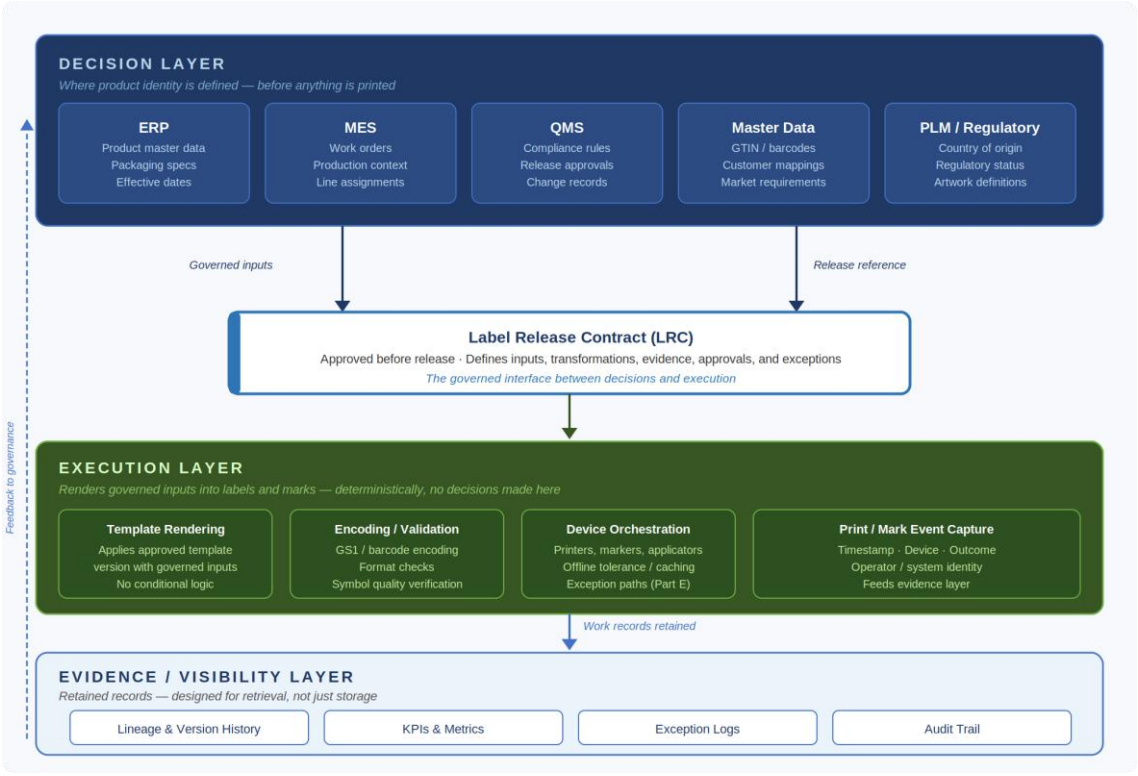

Figure 2: Decision Layer, Label Release Contract, and Execution Layer at Pack-Out

The Decision Layer - everything upstream (ERP, MES, QMS, master data, PLM) - is where product identity is defined: which variant ships to which customer, what the compliance declaration says, when the new packaging specification takes effect. These decisions carry regulatory and commercial risk and require the right people to approve them.

The Execution Layer - the label and mark system - has a precisely defined job: receive governed inputs from upstream, apply approved templates, use validated encoding rules, and capture a record of what was produced. One rule applies throughout: no decisions are made here. If a decision has not been made upstream, execution stops - it does not improvise.

The Label Release Contract (LRC) sits between them. Approved before any label version goes live, it answers five questions explicitly:

- What data drives this label, and where does it come from? (Part A - Governed inputs)
- What is the execution layer allowed to do with that data - and what is it not allowed to do? (Part B - Transformation rules)
- What evidence must be captured to prove what shipped? (Part C - Evidence design)
- Who must approve this release, and what artifacts must be included? (Part D - Release gate)
- What are the authorized procedures for when normal flow breaks? (Part E - Exception paths)

The audit test:

Ask this question: if a customer calls tomorrow and asks which label version was on batch 2024-04-17B, how long does it take to answer? If the honest answer is "a few days, we would have to piece it together" - that is the problem this methodology solves.

Conclusion

Labeling looks simple until it goes wrong at scale. When managed informally, it predictably accumulates drift, variation, and evidence gaps - particularly as the number of products, markets, and sites grows, and as change velocity increases.

The Label Release Contract is the practical tool that changes this. It is not a system, not a software platform, and not a bureaucratic overhead. It is a one-page discipline: a set of questions that must be answered before any label version goes live, and a record of the answers that can be retrieved when they are needed.

The intended outcome is not better templates. It is a labeling capability that is measurable, governed, and repeatable - one that holds up under scale, change, and audit.

→ *Part 2 - Label Release Contract Reference Guide - contains the full LRC template, the detailed rationale for Parts C, D, and E, and guidance on rollout mechanics and operational metrics.*

→ *Part 3 - Completed LRC: Worked Example - shows a fully filled-in contract for a real product launch scenario. Both are designed for practitioners implementing the methodology.*

The author thanks Gerard Ipskamp for peer review and editorial feedback.

Contact Author:

Email: jan.uhrinovsky@gmail.com

LinkedIn: Jan Uhrinovsky | LinkedIn

Website: https://www.yvexa.com

Contact MOMi:

Enquiries: https://www.mom-institute.org/contact-us/

Website: https://www.mom-institute.org/

Email: Info@mom-institute.org