

Where multi-site inventory rollouts actually fail

A teardown of the goods-in → stock-count loop, from someone who has run it on both sides · David Martins

Software companies tend to explain failed rollouts in terms of the customer: they didn't engage, they didn't train, the GM left. Operators explain them in terms of the software: it's too slow, it doesn't match how we work. Both are usually describing the same six failure points from opposite ends. Here they are, with what to check and what to do about each. Nothing here is theoretical — every one of these has cost me money as an operator or cost a client a renewal as a supplier.

THE LOOP, AND WHERE IT BREAKS

A multi-site inventory system only ever tells the truth if a five-step loop closes every single week, in every single site: **order** → **receive** → **record** → **consume** → **count**. Six things break it. In my experience they break in this order of frequency, which is not the order most implementation plans address them.

1. Goods-in is a person, not a process

Symptom: variance looks random rather than trending. Some weeks are perfect, some are catastrophic, and nobody can explain why.

What is actually happening: deliveries arrive between 6am and 8am, into a kitchen mid-prep, and whoever is nearest the door signs the note. On a good day that's the sous chef who checks weights. On a bad day it's a KP who signs for 20kg and 14kg arrives. The system is then correct about what was ordered and wrong about what exists.

- **Check:** pull the delivery notes for one site for two weeks and compare signature names. If there are more than three, the process has no owner.
- **Check:** the ratio of credit notes raised to deliveries received. Under 2% in a fresh-produce operation is not excellence, it is under-checking.
- **Fix:** one named receiver per site per shift, with a deputy, and a physical scale at the door. This is an operational change, not a configuration change — but it is the highest-value thing a rollout can achieve.

2. The unit of measure fights the real world

Symptom: a handful of lines carry impossible variance — 400% one week, negative the next.

What is actually happening: the supplier sells in a case, the kitchen uses grams, the recipe is written in portions and the count is done in "half a box". Each conversion is a chance to be wrong, and someone has usually keyed one of them from memory.

- **Check:** sort variance by value, then check whether the top ten lines share a conversion factor. They usually do.
- **Fix:** for the top 20 lines by spend, agree one unit and enforce it end to end — purchase, recipe, count. Do not attempt this for all 800 lines at go-live; you will stall the project.

3. Transfers between sites are invisible

Symptom: one site is persistently over on a product, its neighbour persistently under, by similar amounts.

What is actually happening: a chef ran out on a Saturday, drove to the sister site and took a case. Operationally this is good management. In the system it is theft at one site and a windfall at the other. In a group with a shared central kitchen this can account for most of the unexplained variance.

- **Check:** correlate variance across sites within the same week. Mirrored numbers are transfers, not losses.
- **Fix:** a transfer must be as fast as taking the case, or it will not be done. If it takes more than about thirty seconds on a phone, the process loses to the Saturday.

4. The count is scheduled against service, not against truth

Symptom: variance improves suspiciously after a manager change, or is far better at one site with no operational reason.

What is actually happening: the count is being done on a Friday afternoon, mid-delivery week, by someone whose bonus depends on the number. Or it is being copied forward from last period for the lines nobody wants to climb a shelf for. A count taken at the wrong moment is not a small error — it invalidates both the period that just closed and the one that just opened.

- **Check:** compare counted quantities line by line across three periods. Identical figures on slow-moving lines mean the shelf was not visited.
- **Check:** the timestamp distribution of the count. A 200-line count completed in eleven minutes did not happen.
- **Fix:** count at close of the last trading day of the period, before any delivery lands. Separate the person who counts from the person judged on the result — this single control is worth more than any report in the product.

5. Wastage has nowhere honest to go

Symptom: wastage is recorded as near zero, and GP is missing anyway.

What is actually happening: the operation has made recording waste feel like a confession. So spoilage, staff food, comps, remakes and the chef's tasting all vanish into variance, where they are indistinguishable from theft. The number is not missing; it is mislabelled.

- **Check:** recorded wastage as a percentage of purchases. Below 0.5% in a fresh operation means it is being hidden, not avoided.
- **Fix:** separate categories for spoilage, staff food, comps and training use, and a fortnight of explicitly not challenging anyone about the numbers. The figure will rise sharply, which is the system starting to work, and needs to be explained to the finance director before it

happens.

6. Nobody has been shown what the report is for

Symptom: full adoption for six weeks, then a slow fade. Logins decline, counts get later, then stop.

What is actually happening: the site is doing three hours of work a week and receiving nothing back it can act on. Head office gets a report; the head chef gets a chore. Every rollout I have seen fail at month three failed here, and it is the one that gets least attention in an implementation plan because it is not a task with a deliverable.

- **Check:** ask a head chef what the last report told them. If they cannot name one decision it changed, adoption is already over — the logins just haven't caught up yet.
- **Fix:** one number per site per week that the chef controls and is beaten on. Not a P&L. Something like top-five-line variance, reviewed on a Monday in two minutes.

THE FIRST 48 HOURS OF ANY ENGAGEMENT

If I am handed a stalled or at-risk multi-site account, this is what I ask for before forming any opinion. It takes about two hours to gather and answers most of the diagnosis.

What I ask for	What it tells me
Login and count-completion data per site, 12 weeks	Whether this is an adoption problem or a product problem. Adoption problems have a decay curve; product problems are flat from the start.
Variance by line by value, top 20, 3 periods	Whether the cause is conversions, transfers, receiving or theft. Each has a distinct signature.
Count timestamps	Whether counting is real.
Delivery notes and credit notes, 2 weeks, one site	Whether goods-in has an owner.
Wastage by category	Whether the operation is being honest with itself.
Names and start dates of site managers	How much of the variance is simply staff turnover.
Fifteen minutes with one head chef, alone	More than all of the above combined.

WHAT THIS MEANS FOR A ROLLOUT PLAN

Four of the six failure points above are operational, not technical. That has a direct consequence for how implementation should be resourced: a consultant who can only configure the product will fix at most two of them, and will be blamed for the other four at renewal.

It is also why I price discovery separately and insist on being on site for at least one delivery and one count. A workshop in a boardroom will tell you what the operation believes it does. Standing at a back door at 6am tells you what it does.

David Martins — 14 years hospitality operations, including a three-venue group and a 3,000-meal-a-day catering operation with 500+ staff; 8 years in hospitality inventory and EPoS software, onboarding 100+ clients and building the onboarding, support and QA functions at an inventory SaaS provider.

+44 7850 450096 · medronho34@gmail.com · Available for implementation, adoption and account-rescue work.