Tracking orders across wholesale, direct and marketplaces without a daily spreadsheet

When orders arrive through four channels in four formats, one order table keyed on a single id is the difference between a system and a daily argument.

GUIDE · 2026-11-03 · 8 MIN READ

A brand selling wholesale, direct and through marketplaces has three order streams that all claim the same inventory. Nobody can see the total, so the number that decides what to buy is assembled by hand each morning from three places, and it is wrong often enough that people stop trusting it.

This is the most common operational failure in consumer brands, and it is not caused by a bad tool. It is caused by having no single id.

What actually goes wrong

Overselling. Direct and wholesale both sold the last twelve units. The wholesale order is the one that gets shorted, because it was entered by hand later.

Chargebacks. Wholesale partners measure your order accuracy. A mis-keyed SKU, a quantity taken from a PDF that changed format, a ship date nobody updated, each one becomes a deduction thirty days later.

Buying blind. Reorder decisions get made from a stock figure that excludes unshipped wholesale orders, so the brand buys what it already has and runs out of what it needs.

The morning ritual. One person exports three files, vlookups them together, and sends a number that expires by lunchtime.

The core move: one order table

Every order, whatever channel it came from, becomes one row in one table:

FIELD	WHY IT EXISTS
Order id, plus the channel's own reference	The join key. Without both you cannot trace back
Channel (wholesale, direct, marketplace)	Everything downstream filters on this
Partner or customer	Who to chase, who to invoice
Order date, requested ship date	The two dates that decide priority
Line items with SKU and quantity	The only level inventory can be allocated at
Status (received, allocated, shipped, invoiced, paid)	The state machine. One field, fixed values
Source document link	The PDF or email the order came from, kept for disputes

Seven columns. Once this exists, the overselling problem becomes a query rather than an argument.

The three rules that make it hold

One status vocabulary, enforced. Not "shipped", "despatched", "sent" and "out the door". Five fixed values, and no row moves backward. Half the reporting problems in a growing brand are a vocabulary problem wearing a database costume.

Allocate inventory at the line level, in one order. Wholesale and direct compete for the same units. Decide the priority rule once, write it down, and apply it in one place. Any rule applied in two places will eventually be applied two different ways.

The source document is evidence. Keep the original PDF or email against the order. When a partner disputes what they ordered, the argument ends in ten seconds instead of a week.

What to automate, in order

1. **Intake.** Wholesale orders are the expensive step because the formats vary. A parser per partner is worth building; a general one is not.
2. **Inventory allocation and the stock figure.** Compute it, never retype it.
3. **Status transitions with a timestamp.** Every change records when it happened, which is what makes the accuracy conversation factual.
4. **Chargeback reconciliation.** Match deductions back to the order that caused them, so you learn which channel or partner is producing them.

Do not automate the buying decision. That is judgement, and a model built on bad data just makes the wrong decision faster.

How to know it is working

- The stock number is the same for everyone, computed once.
- A wholesale order cannot be entered without a partner reference.
- A chargeback can be traced to its order in under a minute.
- You can answer "what is shipping this week across all channels" without opening a spreadsheet.

> *The spreadsheet is a symptom. The absence of a shared order id is the cause, and it costs nothing to fix except the decision to do it.*