Zon Spark

FOR ECOMMERCE OPERATORS · AMAZON · TIKTOK SHOP · DTC

The Ops Leak Kit

Four self-audits for finding where your operation leaks money: the reimbursement worksheet, the Amazon AI agent policy self-audit, the per-SKU profit truth worksheet, and the five loops worksheet. Built to be run in an afternoon, with a pen and your own reports.

2025-26 changed the math on manual ops.

Three policy shifts hit ecommerce operators in quick succession: reimbursements moved to a manufacturing-cost basis with claim windows around 60 days, fee mechanics started punishing both under-stocking and over-stocking, and Amazon's AI agent policy made duct-tape automation a compliance risk. Most brands are still running the old playbook against the new rules. That gap is the leak.

> *"I just want a spreadsheet that updates itself when Amazon changes the price or fees."*
>
> An Amazon seller, in a public Seller Central thread

This kit will not build that spreadsheet for you. It will tell you, in numbers, where your operation is leaking and how big each leak is, which is the part most brands skip before buying tools or hiring help.

WHAT IS INSIDE

Reimbursement self-audit.
What actually changed in the rules, an eight-item documentation-speed checklist, and a quick leak estimate you can run from one revenue number.

AI agent policy self-audit.
The three things the policy requires, eight pass-or-fail checks against your stack, and a one-page automation register to fill in.

Per-SKU profit truth.
The eight cost lines that most per-SKU profit calculations miss, and a worksheet to run your top three SKUs through all of them.

The five loops.
A method for finding your most expensive recurring manual work and putting a dollars-per-month figure on each loop, the same exercise an operations audit starts with.

HOW TO RUN IT

Block an afternoon. You will need access to Seller Central reports (or your channel equivalents), your sourcing costs, and a rough idea of who on your team does what. Every number you write in this kit should come from your own reports, not from memory. Where this kit uses industry figures, they are labeled as estimates and the basis is stated.

WHO MADE IT

Troy Johnston runs his own software company on 130+ automated systems with zero employees, and has been the founder in two multi-million dollar exits, one of them an Amazon software business. Everything in this kit works without buying anything from us. The last page explains what we sell, once, quietly.

Reimbursements: what actually changed

The way Amazon pays you back for lost and damaged inventory changed twice in quick succession. Most sellers are still operating as if neither change happened, and the result is the same in every case: recoverable money expires quietly, sixty days at a time.

1. PAYOUTS MOVED FROM SALE PRICE TO MANUFACTURING COST

In 2025, reimbursements for inventory lost or damaged before a customer order moved to a manufacturing-cost basis: you get paid what the product cost to source, not what it would have sold for. Two details matter more than the headline:

Amazon estimates your cost for you unless you tell it otherwise.
The estimate is built from comparable products and Amazon's own data, and sellers routinely find it lands below their real landed cost.

You can override the estimate.
Provide your own per-SKU sourcing cost documentation. If you do not, you are accepting a discount on every future reimbursement by default. This is the single highest-leverage move under the new policy, and it is one-time setup plus maintenance, not ongoing labor.

2. THE WINDOWS COLLAPSED

Claim eligibility windows that used to allow lookbacks stretching up to 18 months were compressed to roughly 60 days for most categories. Some claim types carry a mandatory waiting period before you may file, then close a few weeks later, leaving a usable filing slot of 15 to 45 days.

OLD REGIME	NEW REGIME
Payout based on expected sale price	Payout based on manufacturing cost (Amazon's estimate unless you document yours)
Lookbacks up to 18 months	Roughly 60-day windows on most categories
File whenever, in bulk	Some claim types effectively open for 15-45 days
Audit quarterly, recover annually	Detect within days or forfeit

The practical consequence: recovery can no longer be batched. A quarterly audit now misses most of what it used to catch. The discrepancy has to be detected close to when it happens, the evidence has to already exist, and the claim has to go in while the window is open.

Why the 20-25% commission services are running the old playbook
Their model was built on mining 18 months of unclaimed history. With 60-day windows the backlog is gone, and recovery at manufacturing cost is a smaller number, so a 25% commission comes off an already smaller payout. The value now is in cost documentation and detection speed, and most commission services were built to mine, not to monitor.

Reimbursement self-audit worksheet

Part A scores your process against what the new rules require. Part B puts a rough annual number on the leak.

A. DOCUMENTATION-SPEED CHECKLIST

- [ ] **Per-SKU manufacturing costs are documented and filed with Amazon.**
  Your real sourced cost is on record, so payouts use your number, not Amazon's estimate.
- [ ] **Cost records are current.**
  When sourcing costs change, the documentation gets updated, with effective dates.
- [ ] **The inventory ledger is reconciled daily, or at worst weekly.**
  Inbound shipments, ledger, customer returns, and removal orders. Quarterly audits miss windows by design.
- [ ] **A missing unit surfaces as a discrepancy within 24-48 hours.**
  Not at the end of the quarter when its window may already be closed.
- [ ] **Evidence is assembled at detection time.**
  Shipment IDs, ledger entries, cost documentation, return scans, pulled together when the discrepancy is flagged, not when someone gets around to filing.
- [ ] **Claims go in days after the event, inside the window.**
  The 60-day window should never be the constraint.
- [ ] **A human reviews every claim before submission.**
  Compliant with how Amazon expects sellers to interact with its systems, and it catches bad claims.
- [ ] **You know what recovery currently costs you.**
  If a service files for you, write the commission here:

  _____
  % of recovered value.

Six or more boxes checked: your process fits the new rules. Fewer: the gap is process, not effort, and it compounds sixty days at a time.

B. QUICK LEAK ESTIMATE

Monthly FBA revenue from your settlement reports	$ _____
Inventory value lost per year, at sale value monthly revenue × 12 × 1.5% (industry estimates put FBA loss and damage at 1% to 3% of FBA revenue; 1.5% is the conservative middle)	$ _____
Recoverable from Amazon per year × your manufacturing cost as a share of sale price (if unknown, use 30%)	$ _____
Likely expiring unclaimed per year × 30% to 60% if recovery runs on a VA and a spreadsheet under 60-day windows	$ _____ to $ _____

Every figure above is an estimate from industry loss-rate ranges, not from your account. Your real number could be lower or higher; the only way to know is to pull your inventory ledger and reconcile it. The expiring portion is the part you control.

Amazon AI agent policy self-audit

Amazon's AI agent policy took effect in March 2026, with enforcement attention ramping since mid-year. If you built automation in 2024 or 2025, the honest starting assumption is that some of it does not comply. Strip away the legal language and the policy requires three things:

SP-API only

Automated tools interact with Amazon through the official Selling Partner API, not by driving Seller Central or the retail site like a human with a browser.

Agents identify themselves

An automated system has to present itself as an automated system. Agents impersonating a human seller session are out.

Human-in-the-loop

Bulk listing edits, price changes, and account-level changes need a human approving them. Agents can prepare and queue; a person clicks the button.

PASS OR FAIL YOUR STACK

Mark each item honestly. Anything in the Fail column that writes changes to your account is your priority list.

1. Every automated write to your account goes through the SP-API. ☐ PASS ☐ FAIL
No unofficial endpoints, no UI-driving write paths of any kind.

2. No browser scripts or RPA bots logging into Seller Central. ☐ PASS ☐ FAIL
Selenium, Playwright, or Puppeteer flows doing listing updates or report downloads fail the SP-API requirement and usually the identification requirement at the same time.

3. No Chrome extensions that write changes. ☐ PASS ☐ FAIL
Tools that bulk-edit listings, adjust prices, or trigger account actions from inside your browser session are acting as you, not as an identified agent.

4. No LLM or computer-use agents driving the Seller Central UI. ☐ PASS ☐ FAIL
Pointing an agent at Seller Central and letting it "handle things" is the most direct possible violation: unidentified agent, UI surface, autonomous consequential actions, all three at once.

5. Repricing and competitive data run on official rails. ☐ PASS ☐ FAIL
No scraping the retail site for competitive data, no price pushes through unofficial endpoints. SP-API pricing feeds only.

6. No shared-login VA-plus-macro stacks. ☐ PASS ☐ FAIL
Offshore VAs running shared logins with macro tools blur the human/agent line. The work is nominally human, but the access pattern looks like a bot because it partly is one.

7. Consequential actions each have a human approval step. ☐ PASS ☐ FAIL
Bulk listing changes, price changes, anything touching account settings. If the answer is "nowhere, it just runs," that is an architecture gap, not a paperwork gap.

8. You can produce a written register of every automation. ☐ PASS ☐ FAIL
Its surface, its identification status, its approval gate. The worst position in an enforcement conversation is not "we had a non-compliant script." It is "we do not know what our automation does."

Your automation register

List every tool, script, extension, VA workflow, and agent that reads from or writes to your Amazon account. Start from the authorized apps list in Seller Central settings, then add what will not appear there: the extension a former employee added, the Zapier flow from 2024, the repricer trial that never got cancelled. Browser-side tooling never shows up in the authorized list, which is exactly the problem.

TOOL / SCRIPT / WORKFLOW	READS OR WRITES?	SURFACE (SP-API OR UI)	IDENTIFIED AS AGENT?	HUMAN APPROVAL GATE?	DECISION (KEEP / REBUILD ON SP-API / KILL)

HOW TO READ YOUR OWN REGISTER

Priority: anything that writes through a UI.
That is the category the policy targets directly. Read-only browser tooling is lower risk, but the trajectory of enforcement says move it.

The identification test.
For each non-SP-API tool, ask whether Amazon can tell the difference between this tool and you. If the answer is no, the tool is impersonating you, and that is the behavior the policy names.

Compliant automation is also better automation.
UI scripts break on every redesign, shared logins are a security hole, and ungated bulk actions are how accounts get damaged by their own tools. The rebuilt version runs faster than the duct-tape it replaces.

The 8 cost lines most brands miss

Ask a founder what the business made last month and you get a number. Ask which of their 40 SKUs actually made money and the answer gets slower, vaguer, and usually wrong. Amazon alone bills through more than 40 distinct fee types; the P&L total hides which SKUs lose money on every unit. Referral and fulfillment fees are the ones everyone counts. These are the eight that quietly break the per-SKU number:

- [ ] **1. Seasonal monthly storage.**
  Storage rates change by season. A SKU that is profitable in March can be margin-negative in Q4 at the same price, and a static model never notices.
- [ ] **2. Aged-inventory surcharges.**
  The clock starts counting at day 181. Slow movers accumulate surcharges that belong on the SKU that earned them, not in an overhead bucket.
- [ ] **3. Low-inventory-level fees.**
  Charged at the FNSKU level. Your margin on a SKU now depends on how well you restock it; under-stocking is billed, not just lost sales.
- [ ] **4. Inbound placement fees.**
  A per-unit cost determined by how you split shipments, invisible if you only look at the fee preview on the listing.
- [ ] **5. Returns and refund administration fees.**
  A SKU with a doubled return rate can look healthy for weeks, because returns come back for weeks. Per-SKU return costs need to land on the SKU.
- [ ] **6. Ad spend attributed per SKU, not averaged.**
  Subtracting total ad spend from total profit is fine for the P&L and useless per SKU. Spend has to be attributed to the products it supported, including halo effects where ads on one ASIN drive sales of another. Average it and your hero SKU subsidizes your losers invisibly.
- [ ] **7. True landed cost, with effective dates.**
  Factory cost, freight, duties, prep, and storage between 3PL and FBA, maintained as data rather than living in old emails and a freight forwarder's invoices. Without it, every downstream calculation is built on a guess.
- [ ] **8. Late adjustments.**
  Reimbursements (now at manufacturing cost), chargebacks, promo clawbacks, and fee corrections trickle in across settlement periods. Any profit number computed the day after a sale is a draft. Models that never reconcile drafts against actual settlement deposits drift, and the drift compounds.

Why the spreadsheet always dies

It encoded the fees and prices of the week it was built, and Amazon changes both continuously. Within a quarter it is quietly wrong, and quietly wrong is worse than absent, because you keep making decisions on it. The fix is not a better spreadsheet; it is a system that re-pulls the inputs and recomputes on events. But before any system, run the worksheet on the next page once by hand. The point is to find out how far off your current number is.

Per-SKU profit truth worksheet

Pick your top three SKUs by revenue and run last month through every line. Pull real figures from settlement reports, the ads console, and your sourcing records. Where you have to guess, mark the cell with a G; the G cells are your data gaps.

LAST MONTH, PER SKU	SKU 1	SKU 2	SKU 3
Units sold			
Revenue			
Referral + fulfillment fees (the ones everyone counts)			
1. Monthly storage (this season's rate)			
2. Aged-inventory surcharges (day 181+)			
3. Low-inventory-level fees			
4. Inbound placement fees			
5. Returns + refund admin fees			
6. Ad spend attributed to this SKU			
7. Landed cost × units (factory, freight, duties, prep)			
8. Late adjustments (reimbursements, chargebacks, clawbacks)			
True margin, $			
True margin, %			

THE RECONCILIATION CHECK

What your model says you earned last month, all SKUs $ _____

What Amazon actually deposited $ _____

Gap $ _____

untraced gaps are how models drift into fiction; if you cannot explain it, your per-SKU numbers inherit it

Decision prompt: any SKU margin-negative after the eight lines? The choices are kill, reprice, or fix the cost line that broke it. Count your G cells too; more than a handful means your real problem is data plumbing, not analysis.

Find your five most expensive loops

A loop is recurring manual work: something a person does weekly or daily that follows roughly the same steps every time. Loops hide because each pass feels small. Priced per month, they are usually the largest controllable cost in an ecommerce back office. Here is how to surface yours.

1. WALK THROUGH LAST FRIDAY

For each person who touches operations, including you, write down what they actually did last Friday. Not the job description, the actual day. Loops live in the gap between those two.

2. NAME THE THREE MOST ANNOYING RECURRING TASKS, IN YOUR WORDS

Annoyance is a surprisingly good proxy for cost. Tasks feel annoying because they recur, resist batching, and interrupt higher-value work.

3. CHECK THE FIVE USUAL SUSPECTS

- [ ] **Reimbursement recovery.**
  Who reconciles the inventory ledger, assembles evidence, and files claims? How many hours, and what commission goes out the door on top?
- [ ] **Profit reporting.**
  Who rebuilds the profit spreadsheet, and how long does "what did we make last month, per SKU" take to answer? When sellers hire developers to build this properly it typically runs $5-15K; most do it with hours instead.
- [ ] **Restock planning.**
  Who turns sales velocity, supplier lead times, freight, and the new fee mechanics into purchase orders, and how often does it go wrong in either direction?
- [ ] **Recurring reporting.**
  Weekly digests, channel summaries, client reports if you are an agency. Count the copy-paste hours per week, per client.
- [ ] **Compliance and account upkeep.**
  Policy changes, flat-file fixes, listing repairs, appeal letters. Irregular but never zero.

VAs and agencies count. If a VA spends 15 hours a week on a loop, the loop costs the VA's loaded rate times those hours, plus the errors and the days of latency. Put it in the table on the next page like everything else.

The five loops worksheet

Price each loop: hours per month × loaded cost per hour. Loaded cost means salary or contractor rate plus everything that rides on it; if you do the loop yourself, use what an hour of your time should be worth, not zero. Be conservative. A conservative number that survives scrutiny beats an impressive one that does not.

LOOP	WHO DOES IT	HOURS / MONTH	LOADED $ / HOUR	$ / MONTH	ERROR / DELAY COST, NOTES
Total					

THEN DO THESE THREE THINGS

- [ ] **Circle the loop you should delete, not automate.**
  Almost every operation has one: a report nobody reads, a reconciliation that duplicates another, a process that exists because of a tool you no longer use. Automating it would preserve the waste. Delete it this week and the worksheet has already paid for the afternoon.
- [ ] **Rank the rest by dollars per month, then check dependencies.**
  The biggest number is not always the first build. A profit model often has to exist before restock or recovery work can be measured properly. Note which loops depend on data another loop produces.
- [ ] **Put the total somewhere you will see it.**
  That monthly figure is what the current way of operating costs. Every month it stands still, it gets paid again.

Everything in this kit works without us.

Run the worksheets, fix what they surface, keep the total where you can see it. If you would rather have the whole exercise done properly, with your real account data instead of estimates, this is what we sell.

AI Operations Audit · $2,500 · 1-2 weeks

This kit, executed against your actual operation: the top five manual loops with dollars per month attached and the basis stated, an automation design for each, and the build order. $7,500 for 8-9 figure brands. The fee is credited in full toward your first system install.

System installs · from $7,500 · fixed scope

The highest-leverage system from the audit, built in production with documentation and handoff. You own everything: no hourly billing, no tool subscription, no lock-in. Built on official APIs with human approval gates, the same architecture the agent policy now requires.

What this is not

Not an "Amazon automation" store scheme; we automate your existing business, never run one for you. You work with the operator who runs these systems in his own company every day.

Start at **zonspark.com/operations**, try the leak estimator at **zonspark.com/estimator**, or write to **hello@zonspark.com**. No newsletter, no drip sequence. Submissions go straight to Troy.

© 2026 Zon Spark. All rights reserved.