

DECISION VELOCITY

The Playbook

How E-Commerce Operators Diagnose and Fix
the Decisions That Take Too Long

By the team at NeonPanel

EXECUTIVE SUMMARY

The Decision Velocity Problem

Every e-commerce business makes the same five decisions repeatedly: whether to replenish a SKU, what the true margin is on a product, whether to adjust ad spend, what the cash position looks like, and whether to run a promotion. These are not complex questions. They are simple questions trapped inside complex data environments.

This playbook exists because of a pattern we have seen across hundreds of operations: the people responsible for strategic decisions spend the vast majority of their time on data assembly — pulling numbers from disconnected systems, reconciling discrepancies, and building a picture that should already exist. By the time the picture is built, the window for action has often closed.

We call this the decision velocity problem. It is not a productivity issue. It is not solved by working harder, hiring more analysts, or buying another dashboard. It is a structural problem — rooted in how data is stored, connected, and made available to the people who need it.

- 01 **A diagnosis.** Based on your audit score, identify which of four operational types describes your business and where your root cause lives.
- 02 **A stress test.** Five decisions that reveal exactly where your operation breaks down — with concrete descriptions of slow vs. fast, and what causes the gap.
- 03 **A fix by bottleneck type.** Not generic advice — specific steps, in order, for each of the four root causes of slow decisions.
- 04 **A working document.** The one-page action plan at the end is designed to be filled in and brought to your team on Monday morning.

This was written by people who have run e-commerce operations for decades and built the platform that solved these problems for our own business. We are not theorising. We are documenting what we learned the hard way — and making it available so you do not have to.

PART 1

What Your Score Means

Your score is between 0 and 100. It is built from four signals: how long your decision cycles take, how many systems you have to pull from, how often bad numbers surface after a decision, and how quickly leadership gets answers to metric questions. Team size affects the dollar model but not the score — a large team can still score 100 with healthy velocity signals.

Question	What it measures	Max points
Q1 — Decision cycle	Time from "we need data" to "decision made"	35
Q2 — Data sources	Number of systems you pull from	25
Q3 — Team size	Scales the \$ model only — not scored	—
Q4 — Post-decision discrepancies	How often bad numbers appear after a decision	25
Q5 — Leadership timeliness	When leadership last got a same-day metric answer	15

YOUR DIAGNOSIS TIER

Three Tiers, One Root Cause

STRONG**80–100**

Velocity signals: Short cycles, few systems, discrepancies caught before decisions are made, timely leadership answers.

Velocity signals are healthy. Your main lever is team size — the larger the team, the more even small amounts of remaining friction compound into real dollar cost. The playbook's fixes will help you maintain this as you scale.

→ [Part 4 — Bottleneck 1 \(maintenance mode\): prevent new manual bridges from forming as the team grows.](#)

MIXED

55–79

Velocity signals: Decisions happen but the path is costly: slower cycles, more sources to pull from, post-decision surprises, or slow leadership answers.

You are paying a friction tax on every decision. The cost is manageable now but multiplies with team size. One or two structural fixes — usually eliminating the most expensive manual bridge or establishing a single source of truth — typically move a Mixed operation to Strong.

→ Part 3 — Bottleneck 1 or 2 are the most common culprits. Part 4 has the specific fix sequence.

AT RISK

0–54

Velocity signals: Long cycles, many systems, frequent post-decision issues, weak same-day leadership clarity.

Decision velocity is a structural drag on the business. At this score, the cost compounds heavily with team size — a 10-person operation at this tier loses the equivalent of months of productive capacity per year. The fix is not faster people. It is a different data architecture.

→ Start with Part 3 — Bottleneck 3 (the number does not exist) or Bottleneck 2 (numbers disagree). These are the root causes for most At Risk scores.

The structural insight

60% of the time your team spends on data assembly is recoverable.

Not through effort. Not through hiring more analysts. Through structure — moving your data to one source of truth so that the time currently spent reconciling is simply no longer required.

The rest of this playbook shows you exactly where the time goes, why it goes there, and how to get it back.

PART 2

The Five Decisions That Reveal Everything

These five decisions expose exactly where an e-commerce operation's velocity breaks down. For each one: the slow version, the fast version, and the gap. Time yourself. Be honest.

DECISION 01

Should We Replenish SKU X?

✗ SLOW**1–3 hours**

Ops lead checks Amazon FBA, the warehouse system, the freight portal, a landed cost spreadsheet, emails the supplier for lead times, checks storage limits, reviews sales velocity. Numbers don't agree. Time spent reconciling before any decision is made.

✓ FAST**5–15 minutes**

Current stock across all locations — FBA, warehouse, in-transit — visible in one view. Landed cost calculated automatically. Lead times stored. The system flags SKUs approaching reorder points. The ops lead reviews and acts.

What causes the gap:

- Inventory data lives in multiple disconnected systems
- Landed cost is estimated or stored in a spreadsheet not linked to inventory
- Lead times are in someone's head or email threads
- No single view combines stock levels, cost data, and demand signals

DECISION 02

What Is Our True Margin on SKU X?

× SLOW**2–6 hours**

Finance estimates an average cost, subtracts marketplace fees (requiring a settlement download and manual breakdown), and arrives at a number they are 70% confident in. If challenged, they cannot drill down to explain exactly where every dollar went.

✓ FAST**Under 1 minute**

Profitability calculated continuously per order using the actual cost of the specific batch consumed (FIFO). Every cost component — purchase price, freight, duties, insurance, fees — allocated at the unit level and fully traceable.

What causes the gap:

- COGS is calculated using averages instead of actual batch costs
- Landed cost components (freight, duties) are not allocated to individual units
- Marketplace fee structures are not systematically categorised
- No system connects the purchase cost of specific inventory to the revenue from specific sales

DECISION 03

Should We Adjust Our Ad Spend on This Product?

× SLOW**30 min + 2–4 week surprise**

Marketing looks at ACoS and ROAS. Numbers look fine. Nobody connects this to the operational reality: actual margin, inventory level, or whether increased demand would cause a stockout. The consequence appears weeks later when inventory runs out or the P&L; reveals the product was sold below true cost.

✓ FAST**15–30 min, no surprises**

Ad spend decisions are informed by real margin data. The person adjusting spend can see current inventory, weeks of cover, true margin after all costs, and break-even ACoS based on actual COGS — not estimated COGS.

What causes the gap:

- Marketing metrics and operational metrics live in completely separate systems
- True margin is not available to the person making the ad spend decision
- Inventory levels and demand forecasts are not visible to the marketing function
- Nobody has defined what "profitable" means at the unit level

DECISION 04

What Is Our Cash Position and Runway?

× SLOW**4–8 hours, stale immediately**

CFO needs: bank balance, outstanding payables (scattered across supplier emails and the accounting system), committed inventory spend (purchase orders that may not be in a system), expected marketplace payouts. The CFO builds a spreadsheet. It takes a day. It is out of date by the time it is finished.

✓ FAST**Under 5 minutes, always current**

Payables, receivables, inventory commitments, and payout schedules all tracked in one system. The cash position is always current. When someone asks "can we afford this PO," the answer is available without building a new model.

What causes the gap:

- Financial data is split between accounting (historical) and operational systems (current)
- Purchase orders are not reflected in financial planning until the invoice arrives
- Marketplace payout timing is unpredictable and not modelled
- Nobody owns the integrated view — finance sees accounting, ops sees inventory

DECISION 05

Should We Run a Promotion This Week?

× SLOW**3–7 days cross-functional**

This decision requires inputs from operations (inventory), finance (margin), marketing (ad performance), and supply chain (lead times). Each input comes from a different person, a different system, and a different meeting. The promotion window may have passed by the time everyone has weighed in.

✓ FAST**One meeting, under 1 hour**

All inputs visible in one place. Inventory levels, margin, ad performance, and supplier lead times all current. The team discusses the strategic question — should we do this? — rather than spending the meeting reconciling whether the numbers even support it.

What causes the gap:

- No shared data layer across operations, finance, and marketing
- Each function prepares its own version of the relevant numbers
- Meeting time is spent aligning on facts rather than debating strategy
- Coordination cost is so high that promotions are often decided unilaterally by one function

DECISION SPEED BENCHMARKS

Where Do You Stand?

Decision	Slow (typical)	Acceptable	Fast (target)
Replenish SKU X?	1–3 hours	30 minutes	Under 15 minutes
True margin on SKU X?	2–6 hours	30 minutes	Under 1 minute
Adjust ad spend?	30 min + 2–4 wk surprise	1 hour	15–30 min, full picture
Cash position?	4–8 hours (stale)	1 hour (weekly)	Under 5 min (live)
Run a promotion?	3–7 days	1 day, one meeting	< 1 hour, same meeting

If most of your decisions fall in the "slow" column, the issue is structural — not about working harder or hiring more people.

PART 3

The Bottleneck Map

Slow decisions are caused by a surprisingly small number of root problems. Identify which ones apply to you.

BOTTLENECK 01**Data Lives in the Wrong Place**

The information exists, but it is in a system the decision-maker cannot easily access, or it requires transformation before it is useful. Your ops lead has to download a CSV from Amazon, reformat it, and paste it into a spreadsheet before anyone can read it.

Common examples:

- Amazon settlement data is only in Seller Central, not in your accounting system
- Warehouse inventory counts are in the WMS but not visible to the finance team
- Freight costs are in emails or PDFs from the forwarder, not in any system at all

Affects most: Type A (Data-Rich, Insight-Poor) and Type D (Numbers Disagree)

BOTTLENECK 02**Numbers Disagree Across Systems**

You have the same metric in two or more places, but the values do not match. Amazon says 340 units. Your warehouse says 312. Nobody is wrong — they are just measuring different things — but nobody has documented or reconciled the differences.

Common examples:

- Amazon FBA count includes reserved units; your internal count does not
- Revenue in Shopify does not match bank deposits after fees, taxes, and refund timing
- COGS in your accounting system uses a different method than your ops team calculates

Affects most: Type D (Numbers Disagree) and Type A (Data-Rich, Insight-Poor)

BOTTLENECK 03**The Number Does Not Exist**

The metric you need has never been calculated. Nobody knows the true landed cost per unit. Nobody has real COGS by SKU — just a blended average or an estimate. SKU-level profitability after all fees and returns does not exist in any system.

Common examples:

- Landed cost is estimated as "purchase price + 20%" rather than calculated from actual freight and duties per batch
- COGS uses average cost instead of actual batch costs (FIFO)
- SKU-level profitability does not exist — only aggregate marketplace-level profitability

Affects most: Type B (Partially Blind)

BOTTLENECK 04**The Person Who Knows Is Unavailable**

The data and systems both exist, but only one or two people know how to extract and interpret the information. When those people are busy, on holiday, or have left the company, decision velocity drops to near zero.

Common examples:

- Only the founder can produce a P&L; that accounts for all cost components
- Only the ops lead knows the reconciliation process between Amazon and the warehouse
- Only the bookkeeper understands how marketplace fees are categorised in the accounting system

Affects most: Type C (Key-Person Dependent)

QUICK IDENTIFICATION EXERCISE

Rank your bottlenecks.

Bottleneck	Applies? (Y/N)	Severity (1–5)	Primary decision it slows
Data lives in the wrong place			
Numbers disagree across systems			
The number does not exist			

The person who knows is unavailable			
--	--	--	--

PART 4

The Fix, by Bottleneck Type

Generic advice is useless. Here are specific steps for each bottleneck type, in the order you should tackle them.

BOTTLENECK 01

Fix for Bottleneck 1: Data Lives in the Wrong Place

The problem is not that data does not exist — it is that extracting it costs more time than analysing it.

STEP 1

Map the actual data journey for your slowest decision

Pick the decision from Part 2 that takes you the longest. Write down every system someone touches to get the answer. This is your data journey map. Most teams have never made this explicit, and seeing it on paper is often the first moment of clarity.

STEP 2

Identify the manual bridges

In that data journey, find every point where a human is acting as the connector — downloading, reformatting, copy-pasting, or verbally relaying information. These are your manual bridges. Each one is a point of delay, a source of error, and a dependency on a specific person.

STEP 3

Replace the most expensive bridge first

Do not try to fix everything at once. Find the single manual bridge that costs the most time or creates the most errors. Usually this means connecting two systems via API, eliminating one system by consolidating, or automating the CSV download-reformat-paste step.

STEP 4

Validate that the bridge stays current

Any connection you build must update in real time or near-real time. A nightly sync is better than a manual process, but it will not get you to the fast benchmark. Aim for systems that push data as it changes, not systems that require someone to pull it.

BOTTLENECK 02

Fix for Bottleneck 2: Numbers Disagree Across Systems

The problem is not that your systems are broken — it is that nobody has established which one is the authority.

STEP 1**List every metric that has more than one source**

Start with the basics: inventory count, revenue, COGS, margin. For each one, write down every system that produces a version of that number. You will likely find 2–4 sources for each metric.

STEP 2**Define the source of truth for each metric**

For each metric, designate one system as authoritative. This does not mean the others are wrong — it means that when they disagree, you know which one to trust. Document this. Share it with your team. This single step eliminates a surprising amount of meeting time.

STEP 3**Understand why they disagree**

The differences usually have a logical explanation: timing differences, scope differences (one includes returns, the other does not), or methodology differences (average cost vs. actual cost). Document the known sources of discrepancy.

STEP 4**Reconcile at a defined frequency**

Set a cadence — weekly for inventory, monthly for financial metrics. Make it someone's explicit responsibility. The goal is not to make numbers identical across all systems but to ensure differences are understood, documented, and within acceptable tolerance.

BOTTLENECK 03

Fix for Bottleneck 3: The Number Does Not Exist

The problem is that you are making decisions without the data those decisions require.

STEP 1**Identify the missing metric**

From Part 2, find the decision where you answered "we don't actually have that number." Name it explicitly: landed cost per unit, COGS by SKU, true margin after all fees, cash position including commitments.

STEP 2**Determine what is needed to calculate it**

Every missing metric is a formula with inputs. Landed cost = purchase price + freight allocation + duties + insurance + other fees, divided by units in the batch. Map the formula. Identify which inputs you have and which are missing.

STEP 3**Close the input gaps**

Missing inputs are usually: data that exists but is not captured systematically, data captured but not at the right granularity, or data that requires a structural change. Start with the first — data that exists but is not captured. Lowest effort, highest impact.

STEP 4**Automate the calculation**

Once the inputs exist, the calculation itself should not require a human. If someone is manually computing landed cost or COGS in a spreadsheet, the number will always be stale, error-prone, and person-dependent. The calculation must happen in a system — automatically, continuously, and traceably.

BOTTLENECK 04

Fix for Bottleneck 4: The Person Who Knows Is Unavailable

The problem is that operational knowledge lives in people, not in systems.

STEP 1**Identify your key-person dependencies**

For each of the five decisions in Part 2, ask: "If the person who usually handles this were unavailable for a week, could someone else produce the same answer?" If no, you have found a key-person dependency.

STEP 2**Document the process, not just the output**

Most key-person dependencies exist because the process was never written down. Interview these people. Have them walk through the process while someone else documents every step, every system touched, every judgment call made.

STEP 3**Reduce judgment calls through system design**

In most documented processes, 80% of steps are mechanical and 20% require genuine judgment. The mechanical steps should be automated. The judgment calls should be documented as decision rules — "if the discrepancy is less than 2%, accept the warehouse number; if greater, investigate."

STEP 4**Build redundancy**

After documentation and automation, ensure at least two people can execute each critical process. This is not about distrusting individuals — it is about ensuring your organisation's decision velocity does not depend on any single person's calendar.

PART 5

What "Good" Actually Looks Like

These benchmarks are drawn from operations that have solved the bottlenecks described in this playbook — businesses that unified their data layer, established sources of truth, and removed the manual bridges.

Replenishment Decision

TARGET**Under 15 minutes from question to decision****What enables this:**

A single view showing current stock across all channels and warehouses, in-transit inventory with expected arrival dates, landed cost from the most recent batch, supplier lead times, and sales velocity with weeks-of-cover projection. The system flags SKUs approaching reorder points automatically.

What it looks like in practice:

NeonPanel calculates dynamic reorder points from lead times, safety stock, seasonality, and sales velocity. The reorder report highlights SKUs at risk. The ops lead reviews the recommendations and acts — instead of spending two hours assembling the data to even begin thinking about the question.

Margin Check by SKU

TARGET**Under 1 minute for any SKU, any period, any marketplace****What enables this:**

True COGS calculated on a FIFO basis from actual batch costs. Every cost component — purchase price, freight, duties, insurance, fees — allocated at the unit level. Marketplace fees automatically categorised. Revenue, cost, and margin computed continuously, not on request.

What it looks like in practice:

NeonPanel's FIFO engine consumes batches in order, assigns the actual cost of each specific batch to each sale, and produces SKU-level profitability that is auditable down to the individual cost component. Over 60 built-in reports. The answer is there before the question is asked.

Financial Close and P&L;

TARGET

Up to 50% faster than manual close processes

What enables this:

Automated journal entries for inventory, COGS, marketplace fees, and payouts — synced to your accounting system in real time. Reconciliation between marketplace data and accounting data happens continuously, so the month-end close is a review, not a reconstruction project.

What it looks like in practice:

NeonPanel pushes invoices, bills, inventory journals, and COGS entries to QuickBooks Online or Xero as transactions occur. P&Ls; and SKU-level profitability breakdowns are available for investor and board reporting without manual spreadsheet assembly.

Cross-Functional Decision

TARGET

One meeting, under 1 hour, with all inputs visible

What enables this:

A unified data layer where inventory, cost, margin, and demand data all live together. When the marketing team asks "can we afford to promote this SKU," the answer is visible in the same system where operations tracks inventory and finance tracks profitability.

What it looks like in practice:

The meeting shifts from "let's first align on the numbers" to "here are the numbers — what should we do?" This does not require new analysis. It requires removing the barrier between the analysis that already exists in different silos.

BENCHMARK COMPARISON

Your Operation vs. the Unified-Data Ideal

This table shows what each audit dimension looks like in a fragmented operation vs. what it looks like when data lives in one source of truth. The gap between the two columns is where your recoverable capacity lives.

Dimension	Fragmented operation (typical)	Unified-data ideal
-----------	--------------------------------	--------------------

Decision cycle	Hours to days — data assembled manually from multiple systems	Minutes — data already unified and current
Data sources touched	3–8+ systems per decision (WMS, Seller Central, spreadsheets, accounting, forwarder portal)	One source of truth — all systems write to the same layer
Discrepancy timing	Post-decision — bad numbers surface after action is taken	Pre-decision — discrepancies flagged automatically before they reach the decision
Leadership metric freshness	Days or weeks — requires someone to build the view on request	Same day, on demand — no assembly required
Recoverable capacity	Baseline: 100% of rework time spent on data assembly	~60% of rework hours recovered through structural unification

The model is illustrative. It does not capture margin protection, speed-to-market advantages, or the cost of burnout in ops teams that spend their days on data reconciliation rather than analysis. The actual value of structural unification is higher than the hours model suggests.

PART 6

One-Page Action Plan

Fill this in. Bring it to your team on Monday.

Our biggest bottleneck decision:

Pick one from Part 2 — the one that takes the longest or causes the most frustration.

How long this decision takes us today:

Target speed (from the benchmarks):

Our primary bottleneck type:

From Part 3: Data in the wrong place / Numbers disagree / Number doesn't exist / Key-person dependent

Secondary bottleneck (if applicable):

The data journey for this decision (systems touched):

#	System / Source	Who does it	Time spent	
1				
2				
3				
4				
5				

The manual bridges (where a human connects two systems):

#	From → To	Could this be automated? (Y/N)
1		
2		
3		

First step to fix it:

Be specific. Not "improve our data" — instead: "enter freight costs into the system within 48 hours of receiving the forwarder invoice."

Who owns this:

By when:

How we will know it worked:

Measure the same decision again in 30 days. Did the time drop?

CLOSING NOTE

Who Wrote This and Why

This playbook was written by people who have lived through every bottleneck described in it.

We have run an Amazon business that hit every wall — fragmented data, estimated COGS, ACoS that told us everything was fine while market share was eroding, an acquisition offer that required clean numbers to evaluate properly. We built NeonPanel because we needed it ourselves, and we spent years refining it against real client operations — not theoretical workflows.

If you recognise your business in these pages, you are not behind. You are where most growing e-commerce operations are. The difference between the businesses that stay there and the ones that break through is usually one structural decision: unifying the data layer so that the humans in the organisation can do what they were hired to do — think strategically, not assemble data.

The data was the bottleneck. It does not have to stay that way.

NeonPanel is the unified inventory and finance platform built for multi-channel e-commerce sellers. We combine inventory management, batch tracking, landed cost allocation, automated COGS, fee reconciliation, and real accounting integration into one platform.

neonpanel.com

We will get sandbox access set up for you this week. Whenever you are ready to start, reach out directly — there is nothing to sign.

[Book a demo →](#)

neonpanel.com