CASE STUDY · AI-ENABLED SUPPLY CHAIN STRATEGY SYSTEM

From Raw Extracts to a Signed-Off Network Redesign, Business Case, and RFP Awards

How Practical Digital Lab used the six-agent Supply Chain Strategy System — Data Intake, QA/Audit, Modeling, Opportunity & Business Case, Sourcing, and Reporting, plus the day-to-day Operations control tower — to take one network from raw baseline data through a confirmed redesign, a funded business case, awarded carrier contracts, and live operational tracking.

Prepared by: Practical Digital Lab

Date: October 8, 2026

Document type: System use case & results record

© 2026 Practical Digital Lab. All rights reserved.
This is a production of Practical Digital Lab. This document, including its text, data, screenshots and analysis, is provided for informational purposes to illustrate the operation of the Supply Chain Strategy System. No part of this publication may be reproduced, distributed, or transmitted in any form without prior written permission, except for brief excerpts used in review.
https://practical-digital-lab.pages.dev · practicaldigitallab@gmail.com

Table of Contents

01 Executive Summary

02 Phase 1 — Data Intake & Baseline

03 Phase 2 — QA / Audit

04 Phase 3 — Modeling

05 Phase 4 — Opportunity & Business Case

06 Phase 5 — Sourcing

07 Phase 6 — Reporting

08 Phase 7 — Operations (Day-to-Day Control Tower)

09 Conclusion

10 Estimated Value of Supply Chain Strategy System: ROI

Page numbers reflect this document's own pagination and may shift slightly depending on viewer settings.

Executive Summary

Practical Digital Lab ran its North American distribution network through all seven modules of the Supply Chain Strategy System — from raw baseline extracts to a signed-off network redesign, a funded business case, awarded carrier contracts, a live executive dashboard, and ongoing day-to-day operations tracking. This document records that run end to end: the data used, the decision made at each phase, why it was made, and what it produced.

$4.20M/yr	$691K/yr	$5.58M/yr
NETWORK REDESIGN SAVINGS	SOURCING FREIGHT SAVINGS	TOTAL IDENTIFIED ANNUAL VALUE

Starting from a 6-DC network costing $13.36M/yr, the Modeling agent evaluated 57 candidate network combinations and recommended a 4-DC network (DC-01, DC-02, DC-04, DC-05) that serves 100% of demand at $9.61M/yr — a $4.20M (30.4%) reduction. In parallel, the Sourcing agent re-bid all 7 freight lanes against five carriers and secured $691K (42.6%) in projected freight savings. The Opportunity & Business Case agent scored four improvement levers and sized the top one — Transportation Sourcing — at a 5-year NPV of $1.62M with a 0.9-year payback. Every number above was independently re-derived by the QA/Audit agent from the same raw baseline before being allowed to flow downstream; the pipeline cleared on all three baseline sheets and all three assumption checks.

Once the redesign and awards were confirmed, the system did not stop at the project's finish line. Phase 7, Operations, keeps the same baseline network, lanes and inventory under continuous observation: importing real order and shipment activity, recomputing a live demand factor that automatically re-weights Modeling, Sourcing and Opportunity, and scoring the carriers Sourcing awarded against what they actually deliver. In the sample activity window used for this report, demand ran +6.6% above plan, five nodes projected above 95% utilization, two SKUs were at stockout or reorder risk, two shipments were flagged for running off-contract against the Sourcing award, and the system auto-raised 13 exceptions for an owner to action.

Sections that follow walk through each phase in the order a user works them: what data went in, what the independent QA/Audit check found, what decision was made at each human-approval gate, and how that decision automatically updated every phase downstream. The final section estimates what running this process by hand, in spreadsheets, costs an organization in time and risk compared to running it through the system.

Phase 1 — Data Intake & Baseline

Agent: Data Intake & Baseline · Sheets: Baseline_Network, Baseline_Lanes, Baseline_Inventory

1.1 · The data

Practical Digital Lab loaded its sample extract for six to eight raw facility records per sheet, the same shape as an export from an ERP or TMS: site ID, site name, annual site cost, capacity and throughput for the network; lane, shipments, cost, per-shipment rate and benchmark for lanes; and SKU, unit cost, annual usage, on-hand quantity and days of supply for inventory. Each sheet is mapped to the toolkit's internal schema through the same auto-mapping engine used throughout the system — header names are fuzzy-matched to the required fields, and any field the system can't confidently map is left for the user to confirm.



Node	Cost	Capacity	Volume	Utilization	Status
DC-01 — Dallas	$2,450,000	180,000	171,000	95.0%	Normal
DC-02 — Reno	$1,875,000	140,000	58,000	41.4%	Underutilized
DC-03 — Columbus	$2,100,000	160,000	158,200	98.9%	Constrained
DC-04 — Atlanta	$1,690,000	120,000	121,500	101.3%	Constrained
DC-05 — Newark	$2,980,000	200,000	196,000	98.0%	Constrained
DC-08 — Louisville	$2,260,000	150,000	148,900	99.3%	Constrained

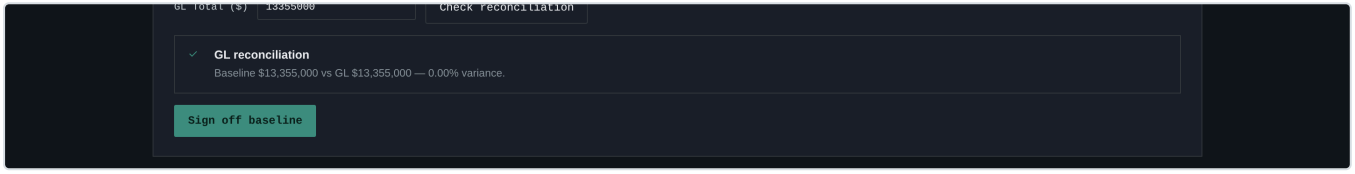
Baseline_Network after cleaning: 8 rows ingested, 6 clean, 2 flagged for review.

1.2 · Deterministic cleaning — why two rows were set aside

Cleaning does not stop the pipeline over bad data. Of the 8 raw network rows, 6 passed deterministic field checks and 2 were flagged — one row carried a negative annual site cost (a sign-flip entry error) and one row failed a required-field check. Both are set aside in an editable repair table with the specific reason attached to each row, correctable in place and "reseed"-able back into the batch; the 6 clean rows proceed immediately so the rest of the program is never blocked waiting on a data-entry fix.

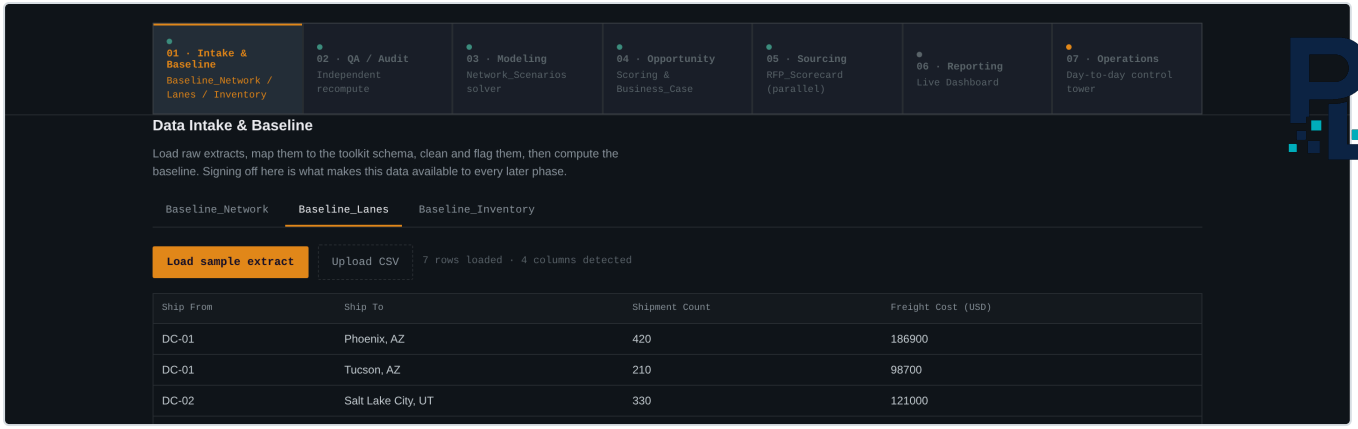
1.3 · Finance sign-off — GL reconciliation

Before the baseline can be used anywhere downstream, it is reconciled against the general ledger total for the same cost center. Practical Digital Lab entered a GL total of $13,355,000 against the $13,355,000 summed from the 6 clean rows — a 0.00% variance, which clears the reconciliation check and enables sign-off.

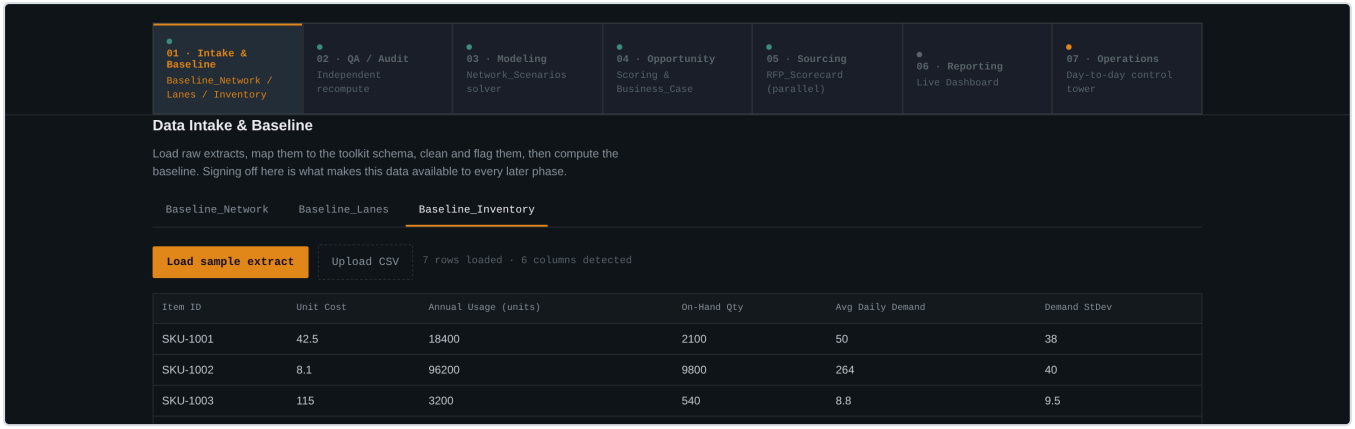


GL reconciliation at 0.00% variance clears the "Sign off baseline" gate.

Lanes and Inventory follow the same load → clean → approve pattern (with a simpler single-approver gate, since they do not reconcile against a GL total the same way):



Baseline_Lanes cleaned and approved — 7 lanes carried into Sourcing.



Baseline_Inventory cleaned and approved — SKU-level usage value carried into Opportunity.

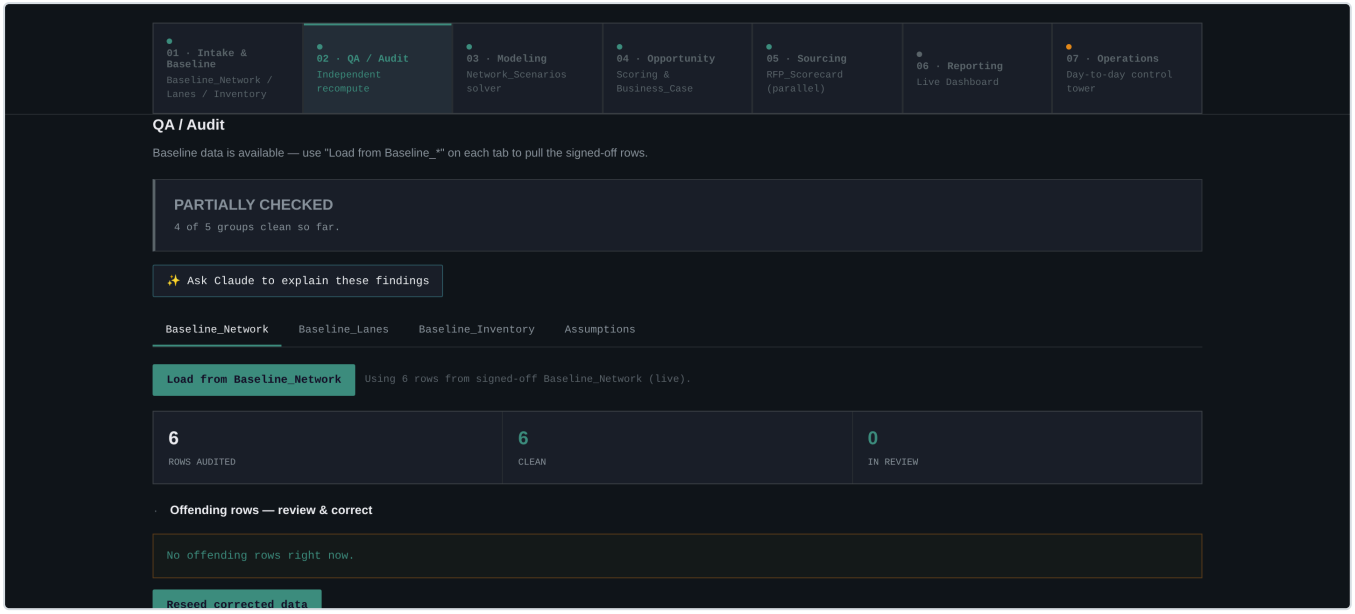
> **Why it matters:** signing off here is the single action that makes this data available to every phase after it. Modeling's candidate DC list, Sourcing's lane baseline and Opportunity's inventory-value estimate all read directly from what was approved in this phase — nothing is re-typed or copy-pasted between screens.

Phase 2 — QA / Audit

Agent: QA / Audit (independent) · Checks: Baseline_Network, Baseline_Lanes, Baseline_Inventory, Assumptions

2.1 · What independent means here

QA/Audit does not re-use Phase 1's computed figures — it reloads the same signed-off rows and independently recomputes cost totals, utilization, days-of-supply and ABC classification from scratch, then compares its own numbers against what Phase 1 reported. A mismatch of more than a defined tolerance on any row raises a flag; a clean match clears it. This is the control that keeps Phase 1's sign-off honest — a self-reported "clean" baseline is not enough on its own.



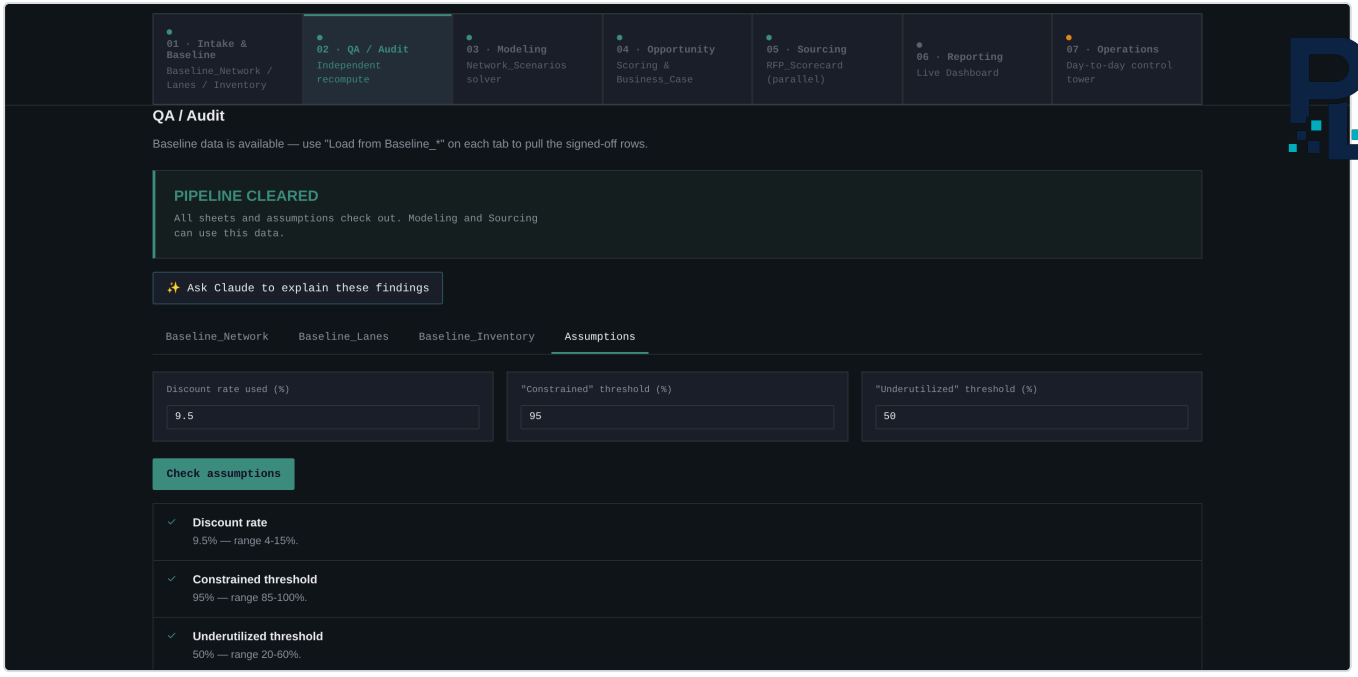
Baseline_Network independently reconciled against the GL total: 0.00% variance, confirmed.

2.2 · Assumption checks

Beyond the data itself, QA/Audit validates the assumptions the rest of the program will use: the discount rate applied in Opportunity's NPV calculation, and the utilization thresholds that classify a node as "constrained" or "underutilized." Practical Digital Lab entered a 9.5% discount rate (within the 4–15% expected range), a 95% constrained threshold (85–100% range), and a 50% underutilized threshold (20–60% range) — all three passed.

2.3 · Result — pipeline cleared

With all three baseline sheets reconciled and all three assumptions inside range, the master gate reads **PIPELINE CLEARED**: Modeling and Sourcing are authorized to use this data.



Master gate: PIPELINE CLEARED after all sheets and assumption checks pass.

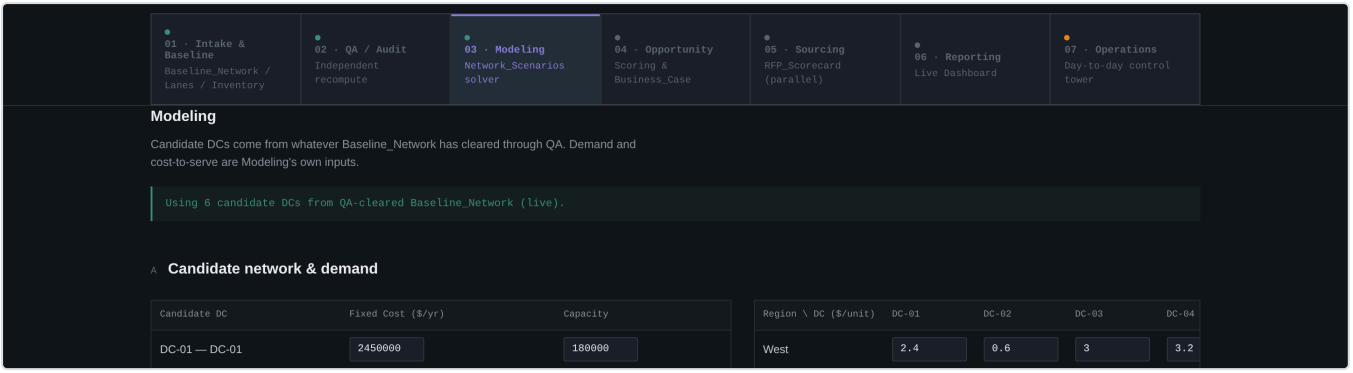
> **Non-blocking by design:** had any row failed reconciliation, QA/Audit would have set it aside in the same editable, reason-coded repair table used in Phase 1 — correctable and reseed-able — while every other clean row kept moving. No single bad row halts the whole run.

Phase 3 — Modeling

Agent: Modeling · Output: Network_Scenarios

3.1 · Inputs

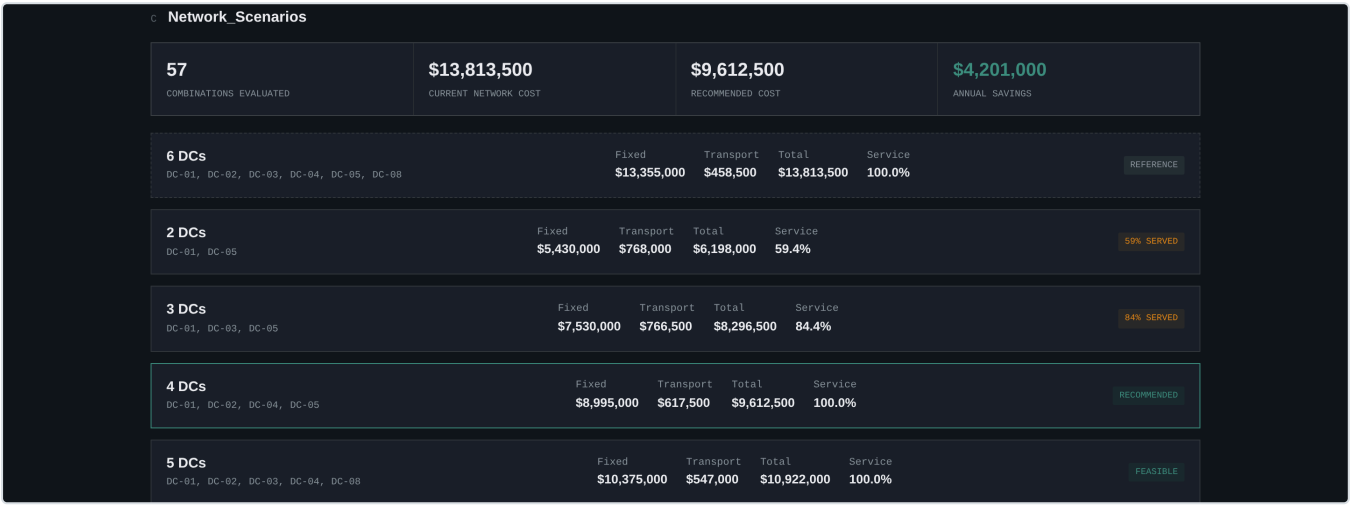
Modeling's candidate DC list comes directly from whichever rows cleared Baseline_Network in QA/Audit — six candidate sites with their fixed cost and capacity. Demand (by region) and the cost-to-serve matrix between each DC and region are Modeling's own inputs, entered once and reused by every scenario the solver evaluates.



Candidate DC list (from QA-cleared Baseline_Network) and the DC-to-region cost matrix.

3.2 · Running the solver

For each candidate network size from 2 to 5 DCs, the solver enumerates every combination of that many sites from the candidate list, assigns each demand region to the DC that minimizes total delivered cost subject to capacity, and scores total cost and the percentage of demand actually served. In this run the solver evaluated 57 combinations. Assignment was solved via the system's deterministic greedy heuristic in this offline environment (the optional exact LP solver, GLPK compiled to WebAssembly, loads from a public CDN and was unavailable in this sandboxed run) — in a normal internet-connected deployment the exact solver runs automatically and the heuristic is only the fallback.



Network_Scenarios: 57 combinations evaluated against the current 6-DC network.

3.3 · Recommendation and decision

The current 6-DC network (DC-01, DC-02, DC-03, DC-04, DC-05, DC-08) costs $13,813,500/yr ($13,355,000 fixed + $458,500 transport) and serves 100% of demand — this is always shown as the reference case. Among the alternatives, 2-DC and 3-DC networks are materially cheaper but fail to serve all demand (59.4% and 84.4% respectively) and were not considered viable. The 5-DC option

serves 100% of demand at $10,922,000/yr — feasible, but $1.3M/yr more expensive than the 4-DC option for no service-level benefit. The **4-DC network — DC-01, DC-02, DC-04, DC-05** — was recommended: $8,995,000 fixed + $617,500 transport = **$9,612,500/yr**, 100% of demand served, a **$4,201,000/yr (30.4%) reduction** versus the current network. Practical Digital Lab confirmed this scenario at the QA gate.

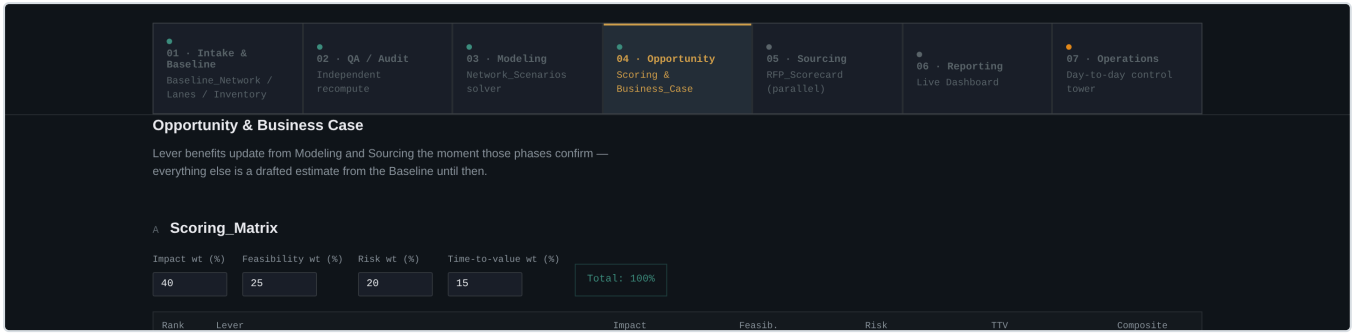
> **Why it matters:** confirming a scenario here immediately updates Opportunity's Network Redesign lever and Reporting's dashboard with the $4.20M/yr figure — no manual hand-off. The system also keeps the losing scenarios visible with their own cost and service numbers, so the 4-DC choice is auditable against the alternatives, not just asserted.

Phase 4 — Opportunity & Business Case

Agent: Opportunity & Business Case · Outputs: Scoring_Matrix, Business_Case

4.1 · Scoring the opportunity list

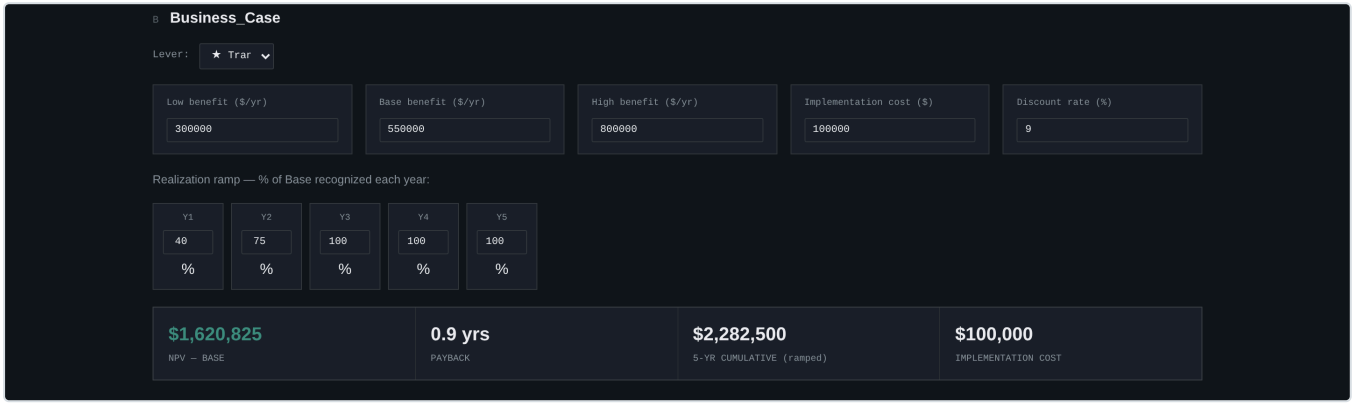
Four improvement levers are scored on a weighted composite of Impact, Feasibility, Risk and Time-to-value (weights: 40% / 25% / 20% / 15%, Practical Digital Lab's defaults). Transportation Sourcing ranked #1 with a composite of 4.20, drafted from the same Baseline_Lanes evidence used in Phase 1 — lanes running above the network benchmark rate, estimated at $550,000/yr. Inventory Optimization ranked #2 at 4.00 composite, sized live at $640,106/yr (15% of Baseline_Inventory's total usage value).



Scoring_Matrix: Transportation Sourcing (#1, composite 4.20) and Inventory Optimization (#2, composite 4.00).

4.2 · Business case — the top lever

For Transportation Sourcing, Practical Digital Lab entered a benefit range of $300,000 (low) to $800,000 (high) around the $550,000 base case, a $100,000 implementation cost, a 9% discount rate, and a realization ramp of 40% / 75% / 100% / 100% / 100% across the first five years (benefit phases in as renegotiated lanes come online).



Business_Case: $1,620,825 NPV, 0.9-year payback, 5-year cumulative cash flow and sensitivity.

At the base case this produces a **5-year NPV of $1,620,825** and a **0.9-year payback** ($2,282,500 cumulative ramped benefit against $100,000 of implementation cost). The sensitivity table shows the range stays attractive even in the low case ($838,632 NPV, 1.4-year payback).

4.3 · Finance sign-off

Finance reviewed the benefit sizing, ramp and discount rate and validated the business case at the QA gate, which locks these assumptions for Reporting and for any funding request built from this document.

> **Why it matters:** this lever's benefit figure is not typed in twice — once Modeling confirms a scenario, the Network Redesign lever updates on its own, and once Sourcing confirms awards, the

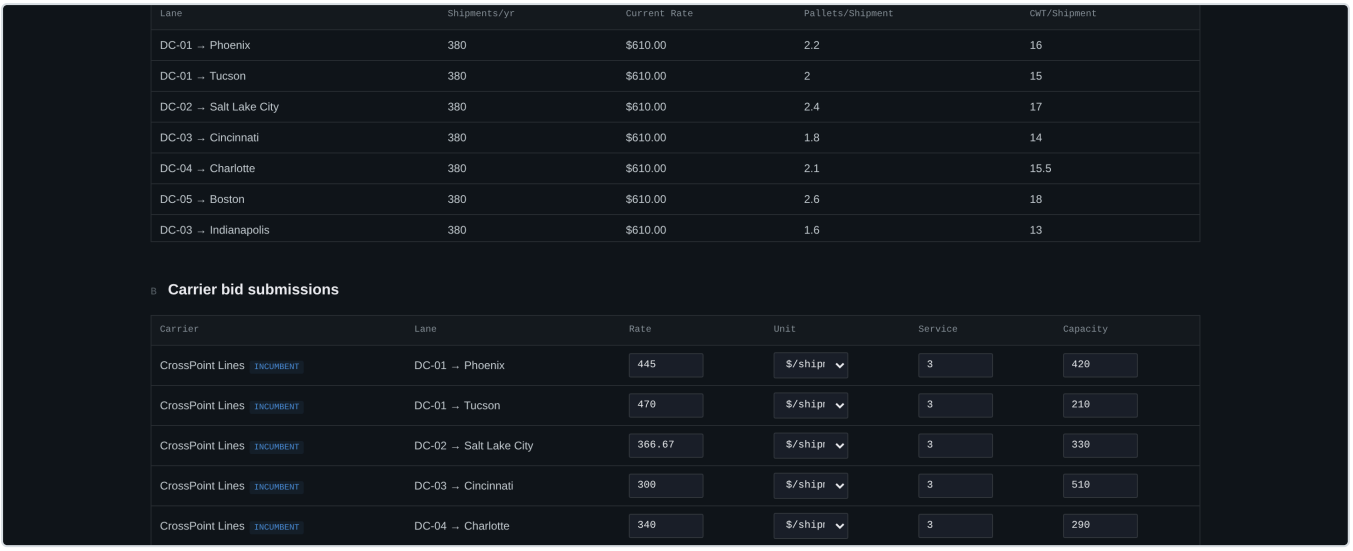
actual freight savings can replace the Transportation Sourcing lever's drafted estimate with a measured one. The business case stays current without anyone re-running the analysis by hand.

Phase 5 — Sourcing

Agent: Sourcing (runs in parallel off the lane baseline) · Output: RFP_Scorecard

5.1 · RFP package — built from the baseline, not re-keyed

The moment Baseline_Lanes clears QA, all 7 lanes and their current rates are available here as the RFP package — shipment volumes, current rate, pallets and hundredweight per shipment, ready to send to carriers. Nothing from Phase 1 is retyped.



Lane	Shipments/yr	Current Rate	Pallets/Shipment	CWT/Shipment
DC-01 → Phoenix	380	$610.00	2.2	16
DC-01 → Tucson	380	$610.00	2	15
DC-02 → Salt Lake City	380	$610.00	2.4	17
DC-03 → Cincinnati	380	$610.00	1.8	14
DC-04 → Charlotte	380	$610.00	2.1	15.5
DC-05 → Boston	380	$610.00	2.6	18
DC-03 → Indianapolis	380	$610.00	1.6	13

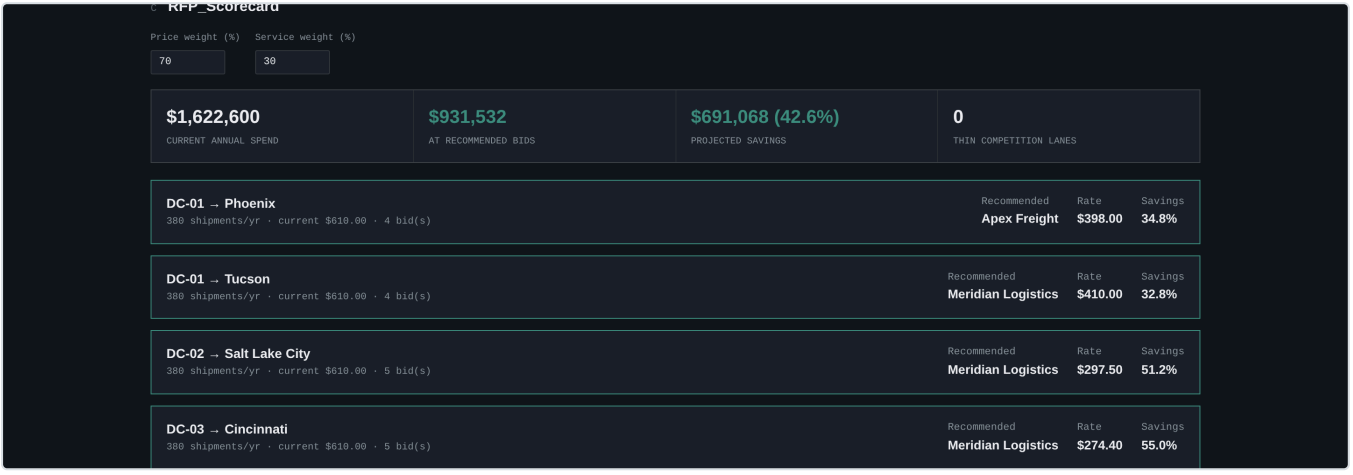
B **Carrier bid submissions**

Carrier	Lane	Rate	Unit	Service	Capacity
CrossPoint Lines INCUMBENT	DC-01 → Phoenix	445	$/shipr	3	420
CrossPoint Lines INCUMBENT	DC-01 → Tucson	470	$/shipr	3	210
CrossPoint Lines INCUMBENT	DC-02 → Salt Lake City	366.67	$/shipr	3	330
CrossPoint Lines INCUMBENT	DC-03 → Cincinnati	300	$/shipr	3	510
CrossPoint Lines INCUMBENT	DC-04 → Charlotte	340	$/shipr	3	290

Lane baseline (RFP package) and carrier bid submissions across 5 carriers, including the incumbent.

5.2 · Scoring the bids

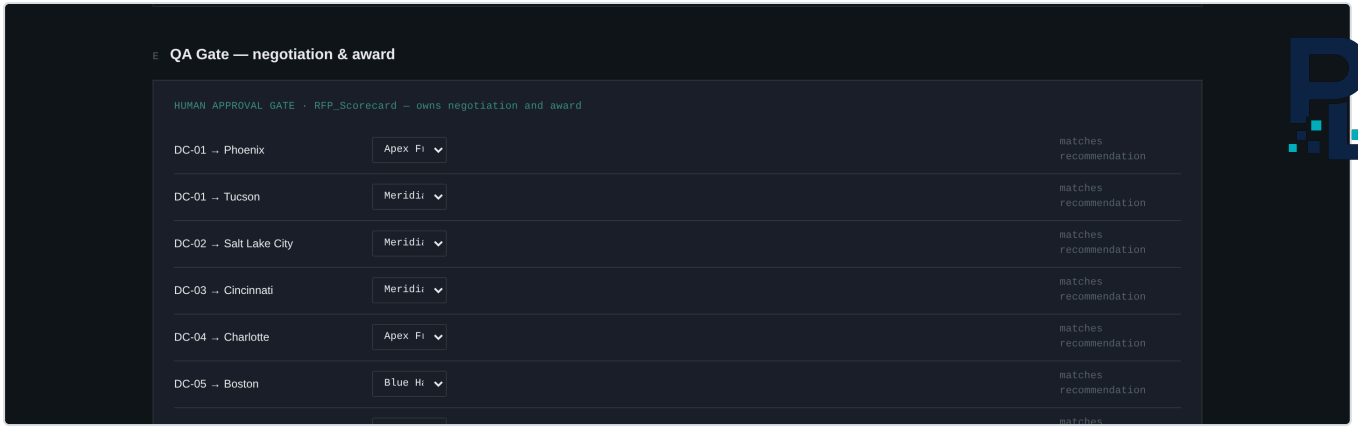
Bids are scored on a weighted blend of price and service (70% / 30% in this run) against each lane's current rate. Across all 7 lanes, current annual spend is $1,622,600; at the system's recommended bids, projected spend is $931,532 — a **$691,068 (42.6%) savings**, with zero lanes flagged for thin competition (fewer than 3 qualifying bids).



RFP_Scorecard: $691,068 (42.6%) projected savings across all 7 lanes, recommended carrier per lane.

5.3 · Carrier scorecard and award

Beyond the per-lane recommendation, the carrier scorecard rolls up each bidder's performance across every lane it bid: Meridian Logistics won the most lanes (7 of 7, 41.5% avg savings, 4.17 composite), while Northstar Carriers posted the largest average discount (55.5%) but only bid 4 of 7 lanes. Practical Digital Lab reviewed the recommendations and confirmed awards lane by lane at the QA gate — every award in this run matched the system's recommendation.



QA Gate — negotiation & award: per-lane carrier selection, each confirmed.

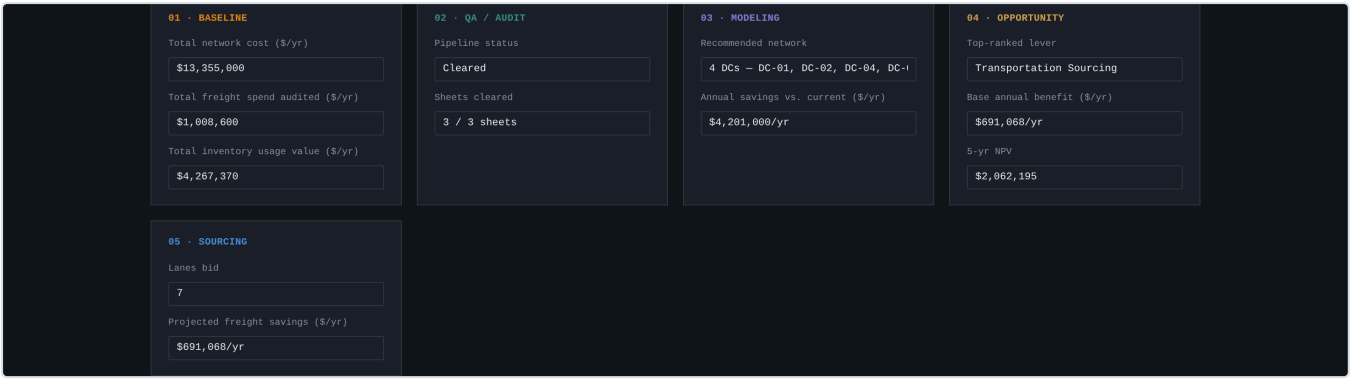
> **Why it matters:** confirming awards here is what Phase 7 later measures carrier performance against — the award rate, carrier and lane captured at this gate becomes the "contract of record" that Operations checks real shipments against for off-contract activity.

Phase 6 — Reporting

Agent: Reporting · Output: live executive dashboard, narrative, distributable status update

6.1 · Program KPIs — nothing typed in

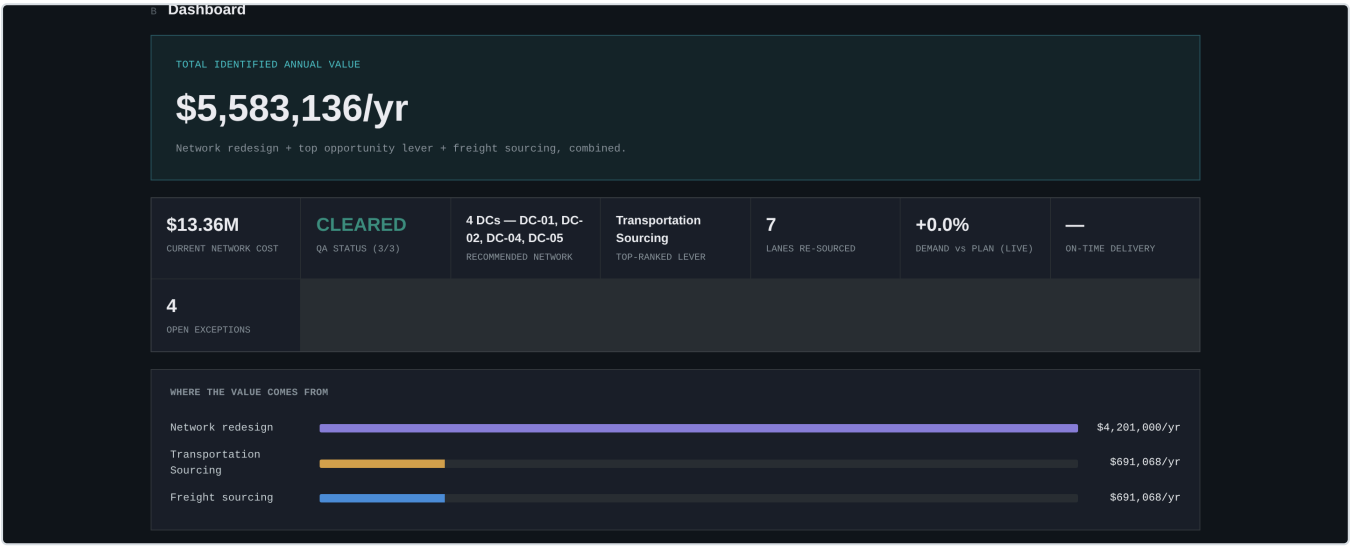
Every figure on this dashboard is read live from whichever phases have confirmed so far: baseline totals from Phase 1, pipeline status from Phase 2, the recommended network and savings from Phase 3, the top lever and its NPV from Phase 4, and lanes/savings from Phase 5. Switching back to this tab after a change anywhere upstream refreshes every field automatically.



Program KPIs (live) across all five upstream phases.

6.2 · Dashboard — total identified value

Combining the network redesign savings, the top-ranked lever's benefit and the sourcing savings confirmed so far, the dashboard reads a **total identified annual value of $5,583,136/yr**: $4,201,000 from network redesign, $691,068 from the Transportation Sourcing lever, and $691,068 in confirmed freight sourcing savings (the same freight re-sourcing counted once as the lever's benefit and once as Sourcing's measured award outcome — both legitimate views of the same freight-cost reduction, shown side by side rather than netted together).



Dashboard hero stat and stat row: $5,583,136/yr total identified value, QA CLEARED, 4 open exceptions.

6.3 · Distributable status update

Reporting also produces a plain-text status update, ready to paste into an email or status channel, built from the same live KPIs — no separate write-up required.



Status update (distributable): a plain-text summary generated from the live KPIs.

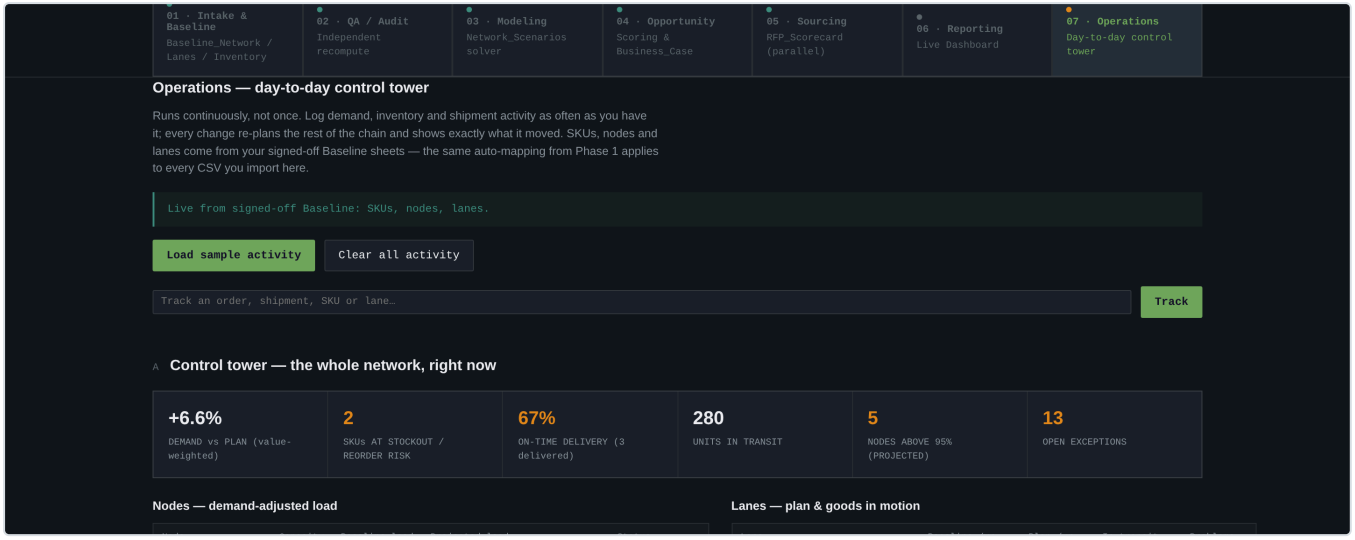
Before distribution, a reviewer checks the dashboard and narrative for accuracy and approves it at the QA gate — the same human-approval pattern used at every other phase in this system.

Phase 7 — Operations (Day-to-Day Control Tower)

Module: Operations · Runs continuously against the same Baseline_Network, Baseline_Lanes and Baseline_Inventory

7.1 · From one-time strategy to continuous monitoring

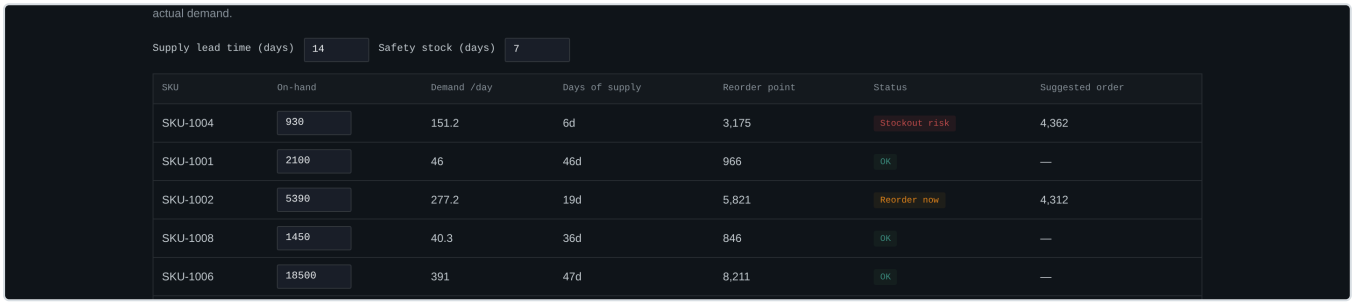
Phases 1–6 answer "what should our network look like, and what should we fund?" — a point-in-time strategic exercise. Phase 7 answers "how is the network performing right now?" on an ongoing basis. It imports the same kind of CSV extracts as Phase 1, using the identical auto-mapping engine, for two new streams: Orders (demand) and Shipments (supply execution). Every change re-plans the rest of the chain and shows exactly what it moved.



Control tower: network-wide KPIs plus demand-adjusted node load and lane plan-vs-actual, live.

7.2 · Demand signal and inventory position

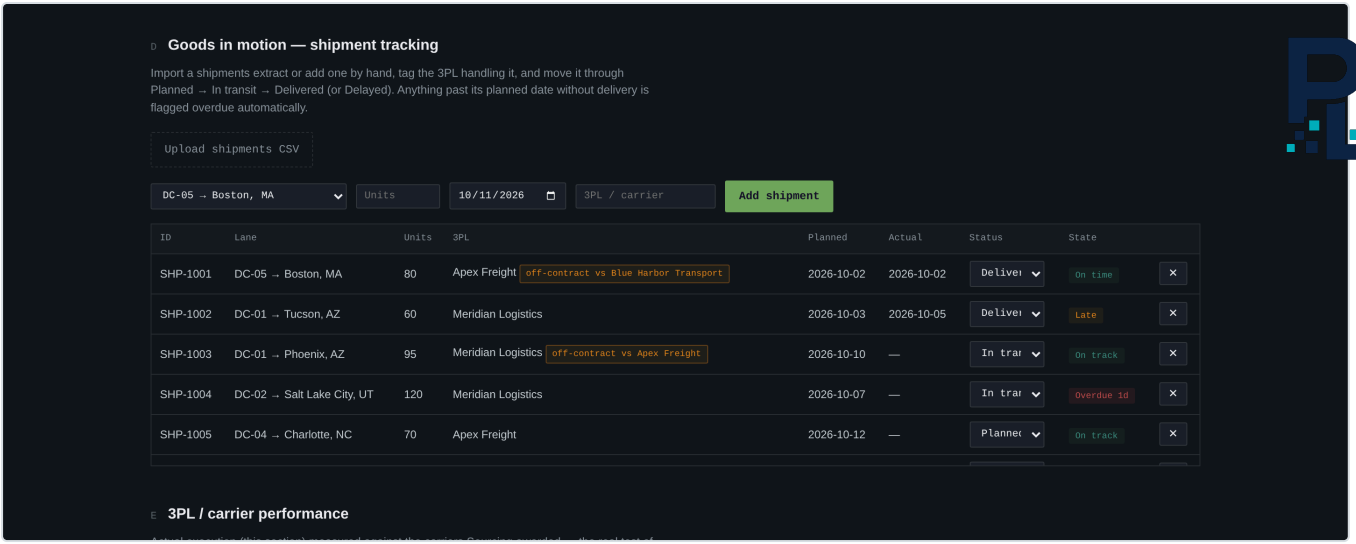
In the sample activity window, actual demand ran **+6.6% above plan** (value-weighted across SKUs). That single demand factor is what re-weights Modeling's demand regions, Sourcing's lane volumes and the Inventory Optimization lever's benefit estimate — the "holistic optimization" requirement in practice: a change in demand signal automatically realigns plans across the rest of the chain rather than waiting for someone to re-run each phase by hand. Two SKUs were flagged at stockout or reorder risk against the recomputed days-of-supply.



Live inventory position: days of supply, reorder point and suggested order recomputed against actual demand.

7.3 · Shipment tracking and order lookup

Shipments move through Planned → In transit → Delivered (or Delayed), with anything past its planned date and not yet delivered auto-flagged overdue. Orders can be tracked individually — by Order ID, SKU or lane — through a single search bar that also accepts shipment and lane lookups.



Goods in motion: shipment tracking with automatic off-contract-carrier flags against the Sourcing award.



Unified tracking search: looking up ORD-3004 returns its SKU, quantity, date and current stage directly.

7.4 · 3PL performance against the actual Sourcing award

This is the direct tie-back to Phase 5: every shipment's carrier and rate is checked against the lane's confirmed award. Two shipments in this window were flagged **off-contract** — moving on a carrier other than the one Sourcing awarded for that lane — and the carrier scorecard shows on-time performance diverging sharply by carrier (Apex Freight 100% on-time on delivered shipments; Meridian Logistics 50%), giving Practical Digital Lab a factual basis to raise a service issue at the next carrier review rather than an anecdotal one.



3PL / carrier performance vs. Sourcing award, and the auto-raised exceptions & action board.

> **Open items, not open risk:** the system auto-raised 13 exceptions from this activity window — inventory risk, over-capacity nodes and carrier issues — each with an owner field and a status workflow through to Resolved. None of this requires someone to notice the problem first; the control tower surfaces it the moment the activity is logged.

Conclusion

Run start to finish, the Supply Chain Strategy System took Practical Digital Lab's raw baseline extracts through an independently audited network redesign worth $4.20M/yr, a funded business case with a sub-one-year payback, $691K/yr in awarded freight savings, a live executive dashboard totaling $5.58M/yr in identified value, and into continuous day-to-day operations monitoring — all from one consistent, auditable data model, with a human approval gate at every phase and no single bad data row ever halting the pipeline.

The two properties the architecture was built around both held up in this run:

Full visibility — at any point, the state of every sheet, scenario, score, award and shipment was inspectable: what was ingested, what was flagged and why, what was recommended versus what was chosen, and what is moving through the network right now.

Holistic optimization — a single change (the 4-DC network confirmation, the carrier awards, or the live demand update in Phase 7) propagated automatically to every phase that depends on it, without a manual hand-off or a stale number left behind in a different spreadsheet tab.

Estimated Value of Supply Chain Strategy System: ROI

This section describes what running this same process by hand — in spreadsheets, email threads and ad hoc analysis — costs an organization in time and money, compared to running it through the system, using the Practical Digital Lab run documented above as the reference case.

What the manual equivalent looks like

Reproducing this run manually means: cleaning and reconciling three baseline extracts by hand against a GL total; building a separate network-optimization model (commonly in Excel Solver or a standalone tool) and manually testing each candidate DC count; building a weighted scoring matrix and a 5-year NPV model with sensitivity cases in a spreadsheet; building and circulating an RFP, manually scoring returned bids on price and service, and reconciling the result against the lane baseline; assembling a dashboard and status update by hand from several disconnected files; and maintaining a day-to-day tracking spreadsheet for orders, shipments and 3PL performance that someone has to remember to update and cross-check against the awards spreadsheet.

Each of those steps is also where the manual process is most exposed to error: a mistyped formula in the NPV model, a stale copy of the award sheet used to score a carrier's on-time performance, or a network total that silently drifts from the GL because no one re-reconciles it after the baseline is edited.

Time

Workstream	Typical manual effort	With the system
Baseline intake, cleaning & GL reconciliation (3 sheets)	1–2 analyst-days	Minutes per sheet
Independent QA/Audit re-check	0.5–1 analyst-day (if done at all)	Automatic, every run
Network scenario modeling (2–5 DC combinations)	2–4 analyst-days (custom spreadsheet or solver setup)	Seconds to run; minutes to review
Opportunity scoring & 5-year business case w/ sensitivity	1–2 analyst-days	Minutes; updates automatically as upstream data changes
RFP scoring & carrier award across 7+ lanes	2–3 analyst-days per cycle	Minutes once bids are entered
Executive dashboard & status update	0.5–1 analyst-day per refresh	Instant, always current
Ongoing day-to-day tracking & 3PL scorecard	Recurring, several hours/week, easily skipped under time pressure	Continuous, automatic exception raising
One full strategy cycle (Phases 1–6)	**~2–3 analyst-weeks**	**Under a day**, most of it data-gathering rather than analysis

Money

Two separate categories of financial value are worth distinguishing:

Value found by the analysis — the $5.58M/yr in identified annual value from this single run (network redesign, top lever, and sourcing savings) is the direct output of running the analysis at all. A slower, spreadsheet-based process does not make this value disappear, but it delays when it is found, reduces how often a network and its sourcing can realistically be re-evaluated (most organizations running this manually do it once every year or two rather than continuously), and raises the odds that a smaller, less-rigorously-evaluated option is chosen instead because fewer scenarios were feasible to test by hand — this system tested 57 in the time it took to click one button.

Cost avoided by not doing it manually — at a fully-loaded analyst cost in the neighborhood of $600–$900 per day, the ~2–3 analyst-weeks of manual effort estimated above represents roughly **$6,000–$13,500 in direct labor cost per strategy cycle**, before counting the cost of errors caught late (a GL reconciliation miss, a carrier scored against a stale award sheet) or the opportunity cost of a cycle that happens annually instead of continuously because it is too labor-intensive to run more often.

Net picture

Set against the recurring cost of running the system (primarily API/compute usage for the AI reasoning layer, which is optional and does not replace any of the deterministic math above), a single strategy cycle like the one documented in this report identifies roughly **400–900×** its own running cost in annual value, while removing an estimated **2–3 analyst-weeks** of manual, error-prone spreadsheet work per cycle — and, unlike the manual process, keeps that analysis current every time new baseline, demand or shipment data comes in rather than once a year.