



VOLUME III · CASE STUDY 001

The Growth Trap.

Preventing margin death-spirals in multi-DC national launches.

EMERGING SPECIALTY PROTEIN · FRESH SAUSAGE
CATEGORY

\$382K

Year-One Cash Burn
Avoided

34.2%

Landed Margin
Maintained

9.6 pts

Trade Spend Reduction

OPERATIONAL INTELLIGENCE · MARGIN ARCHITECTURE · EXECUTION
INFRASTRUCTURE

PRO STRATEGIC SUITE · EDITION 2026

PORTFOLIO & CLIENT POSITIONING USE · CONFIDENTIAL

0 0 | CONTENTS

Inside this document.

A complete diagnostic record of the engagement: findings, root cause, intervention sequencing, results, and operator insight. Built for portfolio presentation, capital conversations, and broker alignment.

0 1	Executive Summary	<i>Engagement framing and headline outcome.</i>
0 2	Performance Snapshot	<i>Before / After / Impact at a glance.</i>
0 3	Detailed Findings	<i>Structured exposure register by operating area.</i>
0 4	Root Cause Analysis	<i>What was actually broken, and why.</i>
0 5	Intervention Strategy	<i>Phased sequencing and milestone gates.</i>
0 6	Results & Operating Impact	<i>Quantified outcomes inside the engagement window.</i>
0 7	Operator Insights	<i>Portable lessons from the engagement.</i>
0 8	Bottom Line	<i>What this case study proves about scaling under PRO.</i>

ENGAGEMENT PARTICULARS

Type: Distribution Autopsy & Staged Expansion Protocol

Timeline: 9-Month Engagement

Prepared by: Scott M. Short, Founder & Managing Principal | PRO (Prime Retail Operations)

Date: April 2026

Authorizations are not a strategy.

A Series A natural-protein brand mistook DC count for momentum. PRO converted a seductive growth accelerator into a disciplined scale engine — before the model lit four distribution centers on fire at once.

A Series A natural fresh-sausage brand had secured KeHE authorizations for four distribution centers on the strength of regional ACV and a high-conviction broker pitch. The founder intended to ship into all four simultaneously and use the launch as the proof point for a Series B raise. The internal narrative was that distribution had been won. The economics had not been audited.

PRO's diagnostic — a forty-seven point Distribution Autopsy run across item setup, deduction history, freight lanes, forecast accuracy, and trade-stack reconstruction — produced a different picture. The current operating model would generate negative contribution margin on every incremental case once the full trade stack and compliance costs were modeled against actual KeHE economics. The brand was, in effect, preparing to pay distributors to take inventory off its hands.

The intervention pivoted the brand from four DCs to a controlled two-DC pilot, rebuilt the landed-margin architecture from the cost line up, capped MCB exposure with hard ROI gates, and installed a protein-specific S&OP cadence between the plant, the 3PL, and KeHE demand planners. Within nine months, fill rate held at 96.8 percent across the pilot, projected year-one cash burn dropped from \$420K to \$38K, and two additional DCs were authorized — but only after the model could earn them.

ENGAGEMENT FRAMING

This was not a turnaround. It was a pre-mortem. The brand had not yet shipped into the additional DCs. The cost of stopping to diagnose was a fraction of the cost of scaling into a margin death-spiral and discovering it on the deduction register six months later.

0 2 | PERFORMANCE SNAPSHOT

At a glance.

The numbers below are the engagement-window deltas — measured against the pre-intervention baseline at scope kickoff. Every line was reconstructed from distributor scorecard data, deduction registers, and the brand’s own P&L.

METRIC	BEFORE	AFTER	IMPACT
Total Trade Spend (% of Gross Sales)	31.4%	21.8%	-9.6 pts
Projected Year-One Cash Burn	\$420K	\$38K	-\$382K
Fill Rate (Pilot DCs)	n/a	96.8%	+96.8 pts
Working Capital Trapped in DC	\$187K	\$61K	-67%
Additional DCs Authorized	0	2	+2 (earned)

\$382K

Year-One Cash Burn
Avoided

34.2%

Landed Margin
Maintained

9.6 pts

Trade Spend Reduction

What the diagnostic actually surfaced.

The findings below are organized as PRO delivered them: as a structured exposure register, not a narrative observation. Every line was quantified before the intervention was scoped.

2.1 Supplier Stage & Production Reality

ISSUE · Plant capacity at 87% utilization with no scenario buffer.

The brand was running approximately \$4.2M in TTM revenue across a natural fresh sausage line with a 35–40 day shelf-life window, co-packed in the Midwest. The co-packer was already running at 87 percent of usable capacity on existing volume. Any incremental KeHE volume would push the plant into make-to-order behavior, with downstream consequences for COGS, freshness at receipt, and on-time-in-full performance.

CONSEQUENCE · Plant utilization above 85 percent eliminates the production buffer that protects MABD windows and shelf-life on receipt. In a 35–40 day shelf-life category, an MABD miss is not a logistical inconvenience — it is a markdown event the brand pays for through fair-share spoilage.

2.2 Trade Stack Blind Spot

ISSUE · Internal model assumed 15% trade. Reality was 27–31%.

The brand's internal trade model assumed 15 percent total trade load. Reconstruction against actual KeHE program economics produced 27–31 percent once MCB pass-through, quick-pay discount, fair-share spoilage allowance, and merchandising fees were included. Two trade categories were entirely unmapped: fair-share spoilage (modeled at 2.8 percent) and merchandising compliance (4.1 percent).

CONSEQUENCE · A trade stack mis-modeled by 12–16 points is not a forecasting error. It is the difference between a positive and negative contribution margin on every case shipped. The brand was capitalizing the launch on numbers that did not exist.

2.3 EDI & Scorecard Compliance

I S S U E · Zero internal ownership of EDI accuracy or scorecard metrics.

EDI 850, 855, 856, and 810 transactions were being handled by the 3PL with no internal owner at the brand. Scorecard GPA was not being read, monitored, or contested. ASN accuracy on the most recent four shipments was 81, 78, 86, and 79 percent — below the threshold at which scorecard demerits begin to compound.

C O N S E Q U E N C E · *Without an internal owner, EDI and labeling errors become silent P&L hits. Compliance is not a back-office function in distributor economics — it is a margin discipline. The absence of an owner was, by itself, a disqualifying signal for further DC expansion.*

2.4 Forecast Accuracy & Inventory Drag

I S S U E · Working capital trapped in DC inventory at 4.5x sustainable load.

Pre-intervention working capital trapped in distributor-side inventory totaled \$187K against an active SKU set that should have supported \$40K–\$60K at the brand's projected turn rate. Forecast accuracy at the SKU-DC level was below 60 percent on the items most exposed to spoilage.

C O N S E Q U E N C E · *Trapped inventory is the most expensive form of capital a brand can hold — it is simultaneously costing carrying expense, generating fair-share spoilage exposure, and preventing the production cadence from staying current. It looks like growth on the receivables side and bleeds margin on the deduction side.*

0 4 | ROOT CAUSE ANALYSIS

Distributor authorizations were being misread as guaranteed sell-through.

The primary failure was a classic founder pattern: equating a signed authorization at a distribution center with a SKU that reorders without subsidy. In every category this is a misread. In fresh protein, where lead times are long and shelf-life requirements at receipt are strict, it is a margin-destroying misread.

The brand lacked a dynamic trade P&L. It treated KeHE as a sales channel rather than a logistics platform with embedded fee schedules. It had no internal owner for EDI, no real-time view of scorecard GPA, and no system for separating real demand from promotional pull-through. The four-DC plan was, in effect, a velocity bet placed against an uncalibrated cost stack.

The deeper diagnosis was that DC count was being measured as a growth metric. In distributor economics, DC count is a fixed-cost multiplier — every additional DC authorization adds compliance load, inventory exposure, freight complexity, and promotional commitment. Without unit economics that survive the trade stack, additional DCs accelerate burn rather than revenue.

THE STRUCTURAL READ

DC count is not a growth metric. It is a fixed-cost multiplier. Without protected unit economics, every authorization compounds the margin liability rather than the revenue.

Diagnose the model before scaling it. Then earn the next DC, do not buy it.

The intervention was sequenced in four phases. Each phase had defined re-test criteria. The brand did not advance to the next phase until the prior one was instrumented and clean.

PHASE I. Distribution Autopsy

Weeks 1 - 2

Forty-seven point audit across item setup, deduction history, freight lanes, forecast accuracy, and EDI error rate. Per-DC risk cards produced. Margin stack reconstructed at the SKU-DC level. Hard-stop register issued: three of four DCs flagged for shipment hold pending architectural correction.

PHASE II. Staged Expansion Protocol

Weeks 3 - 6

Two DCs green-lit on the basis of strongest unit economics and lowest compliance friction. Hard targets installed: 98 percent fill rate floor, 4.2x turn rate minimum, fair-share spoilage capped at 2.8 percent. The other two DCs paused — not abandoned — pending pilot performance.

PHASE III. Landed Margin Architecture

Weeks 4 - 10

List price repriced to absorb realistic cash discounts and quick-pay terms. MCB exposure capped at 4 percent with required incremental ROI gate of 18 percent before any approved event. Fair-share spoilage modeled at 2.8 percent and reserved against. Merchandising spend tied to door productivity rather than blanket calendar coverage.

PHASE IV. Protein S&OP Cadence

Weeks 6 - 36 (ongoing)

Weekly alignment installed between co-packer, 3PL, KeHE demand planners, and the founder. Production cadence tied to receipt MABD windows rather than plant convenience. ASN accuracy escalated to 96 percent in eight weeks. Internal EDI ownership transferred from 3PL to a named role inside the brand.

0 6 | RESULTS & OPERATING IMPACT

Two DCs earned. Two DCs avoided. Margin protected.

Within the nine-month engagement window, the brand moved from a pre-launch model that would have generated negative contribution margin on every incremental case to a two-DC pilot operating with mature CPG discipline. Total trade spend dropped from a projected 31.4 percent to an actual 21.8 percent. Fill rate held at 96.8 percent. Working capital trapped in DC inventory dropped 67 percent.

More importantly, the brand earned two additional DC authorizations on the basis of demonstrated economics — not on the basis of broker enthusiasm or ACV proximity. Those authorizations were extended after KeHE demand planning saw twelve weeks of clean scorecard performance and tight forecast accuracy. The expansion now compounds rather than dilutes.

The avoided cost is the harder number to communicate but the more important one. Year-one cash burn was projected at \$420K under the original four-DC plan. Actual cash burn under the staged-expansion model was \$38K. The \$382K delta is the cost of stopping to diagnose before scaling — the wedge return on a Distribution Autopsy.

M E T R I C	B E F O R E	A F T E R	I M P A C T
Total Trade Spend (% of Gross Sales)	31.4%	21.8%	-9.6 pts
Projected Year-One Cash Burn	\$420K	\$38K	-\$382K
Fill Rate (Pilot DCs)	n/a	96.8%	+96.8 pts
Working Capital Trapped in DC	\$187K	\$61K	-67%
Additional DCs Authorized	0	2	+2 (earned)

What this engagement confirmed.

Portable lessons from the engagement — the kind of diagnostic principles that compound across suppliers, categories, and capital deployments. These are the patterns PRO carries into the next engagement.

- 01** DC count is a fixed-cost multiplier, not a growth metric. Every authorization should be earned by demonstrated unit economics, not bought with broker enthusiasm.
- 02** Trade stack mis-modeled by 10+ points is not a forecasting error. It is the difference between a positive and negative contribution margin on every case.
- 03** EDI and scorecard ownership belongs inside the brand. When it lives at the 3PL, accountability dilutes and compliance becomes a downstream P&L hit.
- 04** Plant utilization above 85 percent eliminates the buffer that protects MABD windows. Production constraint is itself a trade-stack expense.
- 05** The cost of a 48-hour Distribution Autopsy is a fraction of the cost of scaling into a model that has not been audited. Diagnose before you ship.

“DC count is not a growth metric — it is a fixed-cost multiplier. We almost scaled into a margin death-spiral because we were measuring success by doors instead of dollars per case.”

FOUNDER / CEO · EMERGING SPECIALTY PROTEIN BRAND

0 8 | BOTTOM LINE

Disciplined scale beats ambitious scale.

PRO converted a seductive growth accelerator into a disciplined scale engine. By prioritizing unit economics over door counts and earning DC expansion on the basis of demonstrated performance, the brand now operates with the maturity of a \$25M+ peer at a fraction of the burn rate.

The Series B that the original four-DC plan was supposed to support is now better-positioned, not worse. Sophisticated capital reads two-DC clean performance as a stronger signal than four-DC blanket expansion. Discipline is fundable. Burn dressed up as growth is not.

THE PRO MANDATE

PRO can scale suppliers without letting distribution destroy them.

Every engagement begins with a diagnosis and ends with a system the operator can run without us. We do not advise on operations — we install them.

ENGAGE

Scott M. Short

Founder & Managing Principal | PRO (Prime Retail Operations)

Prime Retail Operations · Margin Architecture · Operating Intelligence