THC CONSULTING

VENTURE ASSESSMENT & DIAGNOSTIC SPRINT

D6 — Black Box Bias Audit

Technique 5: Optimism Quantification

NorthBridge Freight Solutions

METRICS AUDITED
6
TAM · CAC · Timeline

PHASE 2 UNIT ECONOMICS
STRUCTURALLY BROKEN
LTV:CAC 0.90:1 optimistic
vs. 3:1 gate required

PHASE 1 INVESTMENT CASE
SURVIVES — CONDITIONAL
Savings real but back-loaded
Payback math tight at 24mo

PREPARED FOR	NorthBridge Freight Solutions · CEO, CFO, VP Business Development
SPRINT DAY	Day 8 — Final Day, Diagnostic Zone
STATUS	Complete — feeds D4 Unit Economics model
DATE	March 2026 · Confidential
PRACTITIONER	Thomas Calvert, Managing Director · THC Consulting

Purpose

Technique 5 applies systematic Bias Adjustment Factors to the VP Business Development's financial projections before those figures enter the D4 Unit Economics model. Every input used in D4 must be bias-adjusted and sourced. The pre-sprint draft financial model is the baseline. VAD standard adjustment factors are applied, with engagement-specific primary research evidence substituted where available — which in this engagement is substantial.

Six metrics are audited: TAM, CAC, time-to-first-revenue, customer ramp and churn, internal savings, and gross margin. The audit closes the Diagnostic Zone. No new data collection occurs after today.

Engagement	NorthBridge Freight Solutions	**Sprint Day**	Day 8 — Final Day, Diagnostic Zone
Practitioner	Thomas Calvert, Managing Director — THC Consulting	**Status**	Complete — feeds D4 Unit Economics model
Date	March 2026	**Classification**	Confidential

1 TAM (Total Addressable Market)

VP BD's Estimate	$500M–$1B

Bias Driver

Category error, not optimism. The VP Business Development used PNW regional freight broker annual revenue as the market size figure. That is the revenue of the buyers, not the technology spend by the buyers. These are not the same number — they differ by two to three orders of magnitude.

Adjustment

Rebuilt bottom-up in D3. 40–60 addressable PNW brokers × GLG-validated willingness-to-pay of $1,000–$1,500/month × 12 months = $345K–$600K annually.

	VP BD (Original)	Bias-Adjusted	Adjustment Factor
TAM	$500M–$1B	$345K–$600K	**~1,000–3,000×**

D4 input: $480K midpoint annual TAM. Use the bottom-up figure exclusively. The VP BD's figure is not referenced in D4.

> **What would justify the original estimate: It cannot be justified. The underlying methodology is wrong regardless of what the inputs are.**

2 CAC (Customer Acquisition Cost)

VP BD's Estimate	Not modeled. S&M; placeholder: $200K (Yr 2) / $450K (Yr 3). Implied CAC: ~$40K/customer (Yr 2), ~$15K/customer (Yr 3).

Bias Drivers — Three Compounding Errors

#	Error	Detail
1	**Error 1 — The placeholder is not a CAC.**	The $200K and $450K budget lines are headcount and marketing estimates. They exclude the fully-loaded cost of a senior sales hire's time per deal, onboarding and implementation support, travel, and extended sales cycle lag before revenue is recognized.
2	**Error 2 — The per-unit math moves in the wrong direction.**	The implied CAC drops from $40K to $15K as customer count scales 7×. This assumes sales efficiency improves dramatically as an unproven team enters an unproven market with an unproven product. In early-stage B2B SaaS, CAC typically rises in Years 2–3 as the easiest prospects are won first.
3	**Error 3 — The competitor-trust barrier is not in the model.**	CU-01 is confirmed negative. The structural trust objection means a portion of the addressable market is unreachable regardless of sales effort. Effective addressable market shrinks; effective CAC rises because remaining qualified sales cycles will be longer and more contested.

Evidence from the Record

> **Sales Lead (8 years experience):** 8–12 months of a senior account manager's time per close. Fully-loaded cost $120K–$150K/year implies $80K–$150K in labor cost per close — before any marketing, onboarding (40–80 hours), or travel. VAD standard bias adjustment (2–3× on client estimate): $40K × 2–3 = $80K–$120K. This is below the labor-only floor, confirming the standard factor is conservative for this engagement.

D4 Inputs — Bias-Adjusted CAC

Scenario	CAC
Optimistic	$100K
Base Case	$150K
Bear Case	$200K

What would justify the original implied CAC: Evidence that a dedicated logistics SaaS hire could close deals in under 3 months with no marketing spend and no extended pilot period. Sales lead testimony directly contradicts this.

3 Time-to-First-Revenue

VP BD's Estimate	~18 months from board approval (assumes 12-month build + Q2 Year 2 sales launch)

Bias Drivers — Four Sequential Delays

	Description
Delay 1	Phase 1A is longer than assumed. The Engineering Lead's honest estimate places data engineering + normalization at 9–15 months. Data normalization is sequential (confirmed Day 8) — it does not run in parallel with platform development.
Delay 2	Phase 2 commercial build is additive, not included. The original model treats Phase 1 completion as the trigger for external SaaS sales. Phase 2 requires additional engineering for external licensing, data partitioning per client, and multi-tenant architecture. Conservative estimate: 6–9 months after Phase 1 deployment.
Delay 3	Sales cycle lag is not modeled. Sales lead estimate: 8–12 months per close from first qualified conversation. The original model assumes customers are available the quarter the product launches.
Delay 4	Sales hire ramp is not modeled. A dedicated SaaS hire requires 2–3 months to recruit, 1–2 months to onboard, and 3–6 months to reach full productivity. The original model shows this hire with no ramp period before pipeline generation.

Cumulative Timeline to First External Revenue

Phase	Duration	Notes
Phase 1A — Data engineering + sequential normalization	9–15 months	Confirmed by Engineering Lead; normalization sequential
Phase 1B — ML build + model training	6–9 months	Requires Phase 1A completion
Phase 1C — Internal deployment + adoption	3–6 months	VP Operations' team must cooperate
Phase 2 — Commercial build (multi-tenant, external API)	6–9 months	Distinct build; not a switch to flip
Sales hire recruit + ramp	5–8 months	Can overlap with Phase 2 build
First qualified sales cycle	8–12 months	Sales lead estimate
TOTAL TO FIRST EXTERNAL REVENUE	**32–48 months**	**From board approval**

> **Original model: ~18 months. Bias-adjusted Base Case: 36–42 months. Standard VAD factor (1.5–2×) applied to the original 18 months yields 27–36 months — and that is before the sequential normalization finding confirmed on Day 8.**

D4 Inputs — Time to First External Revenue

Scenario	Time to First External Revenue
Optimistic	30 months from board approval
Base Case	38 months from board approval

Scenario	Time to First External Revenue
Bear Case	48 months from board approval

What would justify the 18-month estimate: Parallel normalization (confirmed sequential), no Phase 2 build requirement, a sales hire already seated with an existing broker pipeline, and sales cycles under 3 months. None of these conditions exist.

4 Customer Ramp and Churn

VP BD's Estimate	0 → 5 → 35 customers over 3 years. Zero churn modeled.

Bias Drivers

The Year 3 ramp is arithmetically implausible. 30 customers in Year 3 from 1–3 sales hires. Sales lead estimate: 2 qualified conversations per 6 months of full effort per senior hire. At 3 fully-ramped hires, approximately 12–18 qualified conversations per year. At a 25–30% close rate (optimistic for B2B SaaS with a structural trust objection): 3–5 new customers per year. The Year 3 target of 30 customers requires approximately 100% close rate on every qualified conversation from a team that has never sold software, in a market with a confirmed competitor-trust barrier.

Zero churn is an aspirational input, not an analytical one. VAD standard: 15–25% annual churn in Base Case. Sales lead's LTV intuition (2–4 years) implies 25–50% annual churn. The competitor-trust barrier, broker acquisition risk, and build-their-own risk all apply upward pressure.

The time-to-revenue finding (Metric 3) makes Year 2 customers arithmetically impossible. If first revenue arrives at Month 36–42, there are no Year 2 customers. The original 5-customer Year 2 target assumes the product is in market by Month 18. It will not be.

Bias-Adjusted Customer Ramp

Year from Approval	VP BD (Original)	Bias-Adjusted Base Case	Notes
Year 1	0	0	Build in progress
Year 2	5	**0**	Product not yet in market — build runs to Month 30–36+
Year 3	35	**0–2**	Commercial launch Month 36–42; first closes possible late Yr 3
Year 4	—	3–6	First full year with product in market and sales team productive

D4 Inputs — Churn

Scenario	Annual Churn Rate
Optimistic	15%
Base Case	20%
Bear Case	25%

5 Internal Savings

VP BD's Estimate	$1,950K–$2,550K annually. Rounded to $2–3M for board presentation.

> **Note:** This is the metric that matters most for the Phase 1 investment case. The CFO's 24-month payback threshold is measured against internal savings — external ARR is held at zero by her own stated standard. Inflated savings figures directly compromise the payback case.

5a — Dispatcher Efficiency Gain ($400K stated)

Basis: 25 dispatchers × 20% time saving × $80K fully-loaded = $400K. Assumes 100% adoption within 6 months. No ramp modeled.

Bias drivers: The VP Operations has not been formally engaged by the initiative. The dispatcher team has not been consulted. CU-04 (dispatcher adoption) is an active Critical Unknown. The VP Operations' interview flags professional resistance as a realistic outcome. When compensation and incentives are not restructured to support a new tool, adoption rates fall significantly. The 20% time-saving assumption is also unvalidated against actual dispatcher workflow data.

5b — Load Margin Improvement ($1,200K stated)

Basis: 0.5% gross margin improvement on the brokerage book. Management judgment only. Not validated against load data.

Bias drivers: This is the largest single savings line and it rests entirely on an unvalidated assumption. The operations lead's avoidable quality issue rate estimate has not been confirmed by a data pull. The 0.5% improvement assumes the AI model immediately performs materially better than current dispatcher judgment — itself untestable until the model is built, deployed, and operating in production for a period sufficient to generate learning.

5c — Carrier Penalty Avoidance ($350K stated)

Basis: Operations lead estimate of current penalty rate. Single internal source.

Bias drivers: This line is more operationally traceable than the others — penalties are recorded. The bias is less about methodology and more about attribution: the platform would need to demonstrate causation, not just correlation, with penalty reduction. The figure is directionally plausible if the system performs as intended.

Internal Savings Summary — Bias-Adjusted

Category	VP BD Annual	Bias-Adjusted Yr 1 (post-deployment)	Bias-Adjusted Yr 2
Dispatcher efficiency gain	$400K	$112K–$140K	$260K–$300K
Load margin improvement	$1,200K	$240K–$360K	$600K–$720K
Carrier penalty avoidance	$350K	$210K–$245K	$280K–$315K
TOTAL	**$1,950K**	**$562K–$745K**	**$1,140K–$1,335K**

> **The original stated figure overstates Year 1 savings by approximately 2.6–3.5×.**

Payback implication

Parameter	Value
Build cost (Engineering Lead revised)	$2.0M–$2.5M
Required annual savings (24-month gate)	$1.0M–$1.25M from deployment Month 1
Bias-adjusted Year 1 savings	$562K–$745K
Year 1 shortfall vs. required run-rate	**$255K–$688K**
Year 2 savings — clears the threshold?	Approaches but does not reliably clear across the full 24-month window

6 Gross Margin on External SaaS

VP BD's Estimate	Not stated. Implied 100% (no cost of revenue line in model).

Bias Driver

The financial model carries all external ARR to the bottom line with no cost of revenue. For a data-intensive B2B SaaS platform serving freight brokers, this understates the cost structure significantly. Each customer requires: implementation and data integration (40–80 hours), ongoing customer success support, ML inference compute costs scaling with usage, and data hosting. Comparable B2B SaaS with meaningful implementation requirements typically runs 55–70% gross margin at scale — lower in early years when per-customer support costs are highest.

Scenario	Gross Margin Applied
Optimistic	68%
Base Case	60%
Bear Case	50%

Synthesized Picture

Phase 2 — External SaaS Unit Economics

At bias-adjusted inputs, the external SaaS business does not produce a viable LTV:CAC ratio under any scenario that uses evidence-grounded assumptions.

	Optimistic	Base Case	Bear Case
ARR per customer (GLG WTP)	$18K/yr	$15K/yr	$12K/yr
Avg. customer life (churn-adjusted)	5 years	4 years	3 years
Gross LTV	$90K	$60K	$36K
CAC	$100K	$150K	$200K
LTV:CAC	**0.90:1**	**0.40:1**	**0.18:1**
vs. 3:1 gate	**FAIL**	**FAIL**	**FAIL**

VAD threshold for a viable external business: LTV:CAC > 3:1 in the Base Case. Base Case result: 0.40:1 — not a marginal miss. A structural failure. Even the Optimistic scenario does not approach 1:1. The economics are compressed from three directions simultaneously: a low WTP ceiling (GLG-validated at $1K–$1.5K/month), a high CAC floor (sales lead validated at $80K–$150K labor cost alone), and a small addressable market ($345K–$600K annual TAM). These three forces are each independently supported by primary evidence and mutually reinforcing. No single fix resolves the unit economics problem.

Phase 1 — Internal Savings Payback

Parameter	Original Model	Bias-Adjusted Base Case
Build cost	$1.5–2M	$2.0–2.5M (Engineering Lead revised)
Required annual savings (24-mo gate)	$750K–$1M+	$1.0M–$1.25M
Year 1 actual savings	$1.95M	**$562K–$745K**
Year 2 actual savings	$1.95M	$1.14M–$1.34M
24-month payback — clears CFO's gate?	Yes (model)	No (Yr 1 short by $255K–$688K)

Phase 1 investment case survives but is narrower and harder than the original model implies. Year 1 savings are substantially back-loaded. The payback math is tight and depends on adoption rate assumptions that are unvalidated and actively contested by the VP of Operations.

Verdict Implication for D4

Phase 2 — External SaaS

Unit economics are structurally broken at current evidence levels. This is not a calibration problem — it is a model design problem. The financial case for Phase 2 cannot be constructed from available evidence in a way that clears any standard investment threshold. T5 makes Phase 2 harder to defend than even the CU-01 competitor-trust finding, because CU-01 is theoretically addressable through structural solutions (spinout, white-label, neutral partnership). The LTV:CAC failure is driven by forces that structural solutions do not fix: a low WTP ceiling and a small addressable market remain regardless of how the entity is housed.

Phase 1 — Internal Platform

A conditional investment case survives. Savings are real but back-loaded. Payback math is tight against the CFO's 24-month threshold and depends on dispatcher adoption assumptions that the VP Operations has not confirmed. The investment decision hinges on whether preconditions — principally the VP Operations' structural engagement and a formal data audit — are met before capital commitment.

> **D4 construction rule: Every input in D4 must show (a) the original figure, (b) the T5 adjustment factor and its source, and (c) the bias-adjusted figure used in the model. The pre-sprint draft figures are not usable as D4 inputs.**

Thomas Calvert · Managing Director, THC Consulting · thomas@sprintverdict.com
NorthBridge Freight Solutions · D6 Black Box Bias Audit · Technique 5 · Day 8 · Confidential