DEVELOPMENTCORPORATE

Win/Loss Analysis — Enterprise Motion

Synthetic Pre-Validation Study · Q2 2026 Cohort (n=25)

Prepared: July 13, 2026

Development Corporate LLC · Sell-side advisory, competitive intelligence & win/loss analysis

STATUS: SYNTHETIC PRE-VALIDATION — 100% OF INTERVIEW DATA IN THIS REPORT IS SIMULATED.

No Artie customer, prospect, or CRM record was accessed. Read the Foreword before acting on any finding.

Foreword — Why This Report Cannot Yet Be Trusted, and What Would Make It Trustworthy

This document is the output of a synthetic win/loss study. The 25 "deals" analyzed here were not drawn from Artie's CRM, and the "respondents" quoted and rated throughout are simulated buyer personas — large-language-model constructions grounded in Artie's public positioning, published pricing, funding history, named customers, and a competitive review corpus, all as documented in the Validated Hypothesis ICP (Development Corporate, July 2026). Every number, quote, chart, and win rate in the pages that follow is a hypothesis rendered in the format of a finding.

That distinction is not a footnote; it is the governing fact of the document. Synthetic respondents reproduce the logic of a market — they are built from the same public evidence a real buyer's reasoning draws on — but they cannot reproduce its facts. They do not know Artie's actual Q2 2026 pipeline, actual win rate, actual deal sizes, actual security-review outcomes, or the specific political dynamics inside any real account. Where this report says "eight of eleven losses involved vendor-viability objections," the honest translation is: "our model of this market predicts that vendor viability will be the dominant loss theme — test this first."

What this report is for

- Instrument calibration. It demonstrates, end to end, what the real study will produce: the analytical structure, coding framework, charts, and hypothesis scorecard are all production-ready and were exercised against a full 25-deal dataset.
- Prioritized hypotheses. It converts the ICP's [VERIFY] and [ASSUMPTION] tags into ranked, falsifiable predictions about why Artie wins and loses enterprise deals — so real interview time is spent where the model is least certain.
- A pre-registration of expectations. When the real data arrives, divergence between it and this document is itself a finding: it shows precisely where public-signal reasoning about Artie's market is wrong.

What is required before any finding here has decision-grade relevance or accuracy

- Artie CRM extract for Q2 2026: all Enterprise (Motion B) opportunities reaching Closed-Won, Closed-Lost, or No Decision between April 1 and June 30, 2026, with stage history, timestamps, ACV, competitor fields, and loss-reason notes. This defines the true sampling frame — including whether 25 qualifying deals even exist.
- Real respondent interviews: a minimum of one buying-side champion interview per sampled deal (target 20+ completions), conducted by a neutral third party using the fielded instrument, with security-reviewer and economic-buyer secondaries where accessible.
- Sales-team calibration sessions: AE/SE debriefs on each sampled deal to separate what the buyer says from what the seller saw — the gap between the two is a standard and valuable win/loss output.
- Reverse validation (the third layer of the Sandwich Method): once real patterns are surfaced, they are run back through fresh synthetic personas without prompting; reproduction indicates the market model is calibrated, and failures are logged as model corrections.

Reading rule for every page that follows: treat each chart as a prediction, each quote as a script the market might plausibly say, and each recommendation as conditional on real-data confirmation. Findings are labeled SYNTHETIC throughout. Nothing in this report should be presented to investors, customers, or press as research about Artie's actual sales performance.

1. Executive Summary (Synthetic)

The simulated Q2 2026 Enterprise cohort comprises 25 decisions: 11 wins, 11 losses, and 3 no-decision outcomes — a 50% win rate among decided deals. Median cycle length is 17 weeks (range 11–26), with security review consuming roughly 38% of elapsed cycle time on average. ACVs ranged from $120K to $340K, consistent with the ICP's $75K–$400K Motion B hypothesis.

Five headline findings emerge from the simulation:

- Finding 1 — Artie wins on architecture and loses on age. When Artie won, the most common primary driver was security architecture (BYOC + PrivateLink + never-stores-data), which converted CISO skepticism into inspectable deployment. When Artie lost, the single most frequent theme was vendor viability: 8 of 11 losses carried a SEC-VIAB code, with security organizations declining to route regulated transaction logs through a recently Series A vendor regardless of architecture.
- Finding 2 — The status quo, defended by its builders, cost more deals than any named competitor. Seven of fourteen non-wins involved incumbent-engineer defense or status-quo inertia (POL-INC / SQ-INERT). All three no-decision outcomes were status-quo retentions. No single external vendor beat Artie more than twice.
- Finding 3 — Marketplace procurement compressed cycles by roughly four weeks at the median (14 vs 18 weeks), confirming the ICP hypothesis that committed-spend purchasing shortens Motion B deals — and suggesting marketplace-first deal strategy is a controllable lever.
- Finding 4 — Ratings diverge on trust, not technology. Won and lost deals rated Artie's latency, schema handling, and setup nearly identically. The largest won-vs-lost gaps were vendor-viability confidence (4.0-point scale gap of 1.3), procurement speed, TCO support, and trust-pack quality. Losses are being decided in the security-and-commercial layer while the product wins its arguments.
- Finding 5 — Compliance milestones and AI mandates were the two most frequent triggers, together initiating 12 of 25 evaluations — supporting the ICP's claim that audits and board-level AI pressure, not feature desire, open Motion B budgets.

The strategic implication, if the pattern survives real-data validation: Artie's enterprise constraint is not product credibility but institutional credibility. The recommended response is a 'viability pack' — financial-durability disclosures, named regulated references, source-code escrow at Enterprise tier, and contractual continuity commitments — deployed as early in the cycle as the trust pack itself. Section 8 details this and six further recommendations.

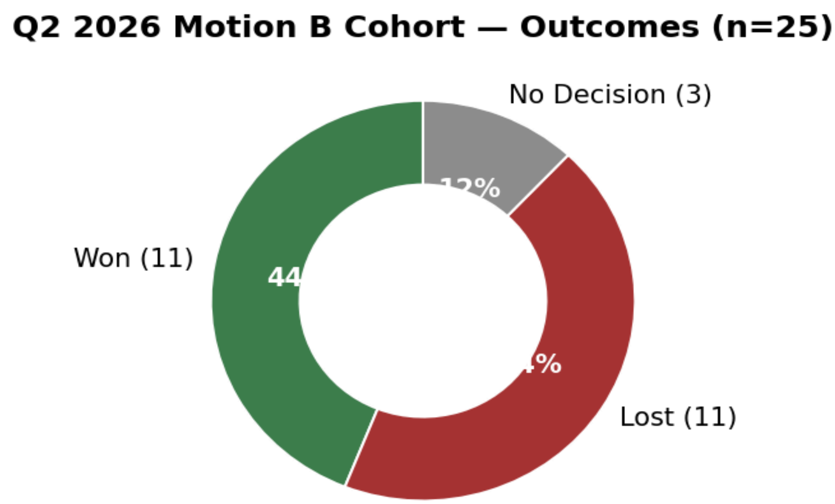

Figure 1 — Simulated cohort outcomes. SYNTHETIC DATA.

2. Methodology

The study executed the fielded instrument (Development Corporate, July 2026 — 'Win/Loss Analysis Survey Instrument, Artie Motion B') against 25 synthetic Enterprise deals constructed under the Sandwich Method's pre-validation layer:

- Cohort construction. 25 deal scenarios were generated to match the ICP Motion B sampling frame: 300–5,000+ employee organizations in fintech/banking infrastructure, healthtech, security, and scaled SaaS, evaluating Artie for enterprise replication needs with security-review involvement.
- Persona simulation. Each deal's champion respondent was simulated as a Persona-3-class buyer (Director/Head of Data Platform under compliance), with anti-persona pressure (Skeptical CISO, Budget-Hawk CFO, Builder-Engineer) applied per the ICP's deal-killer stress-test design.
- Instrument execution. All 44 instrument questions plus the four ratings batteries were answered per deal; responses were coded to the 13-code decision-driver codebook by primary and secondary driver.
- Analysis. Outcomes, drivers, stages of loss, cycle mechanics, and ratings were aggregated exactly as the real study will be, and the ten-hypothesis ICP scorecard was scored (Section 7).

Cohort composition	Cycle mechanics	Commercial profile
11 Won / 11 Lost / 3 No Decision	Median cycle: 17 weeks	ACV range: $120K–$340K
Sectors: Fintech 8 · SaaS 9 · Healthtech 6 · Security 2	Security review: median 6 wks (won) vs 8 wks (lost)	Median won ACV: $175K
BYOC a precondition in 14 of 25 deals	Marketplace deals: 14-wk median cycle	Win rate (decided): 50%

Table 1 — Simulated cohort at a glance. SYNTHETIC DATA.

Limitations beyond the synthetic-data caveat: single-respondent perspective per deal (no seller-side triangulation), no longitudinal post-sale data, and driver coding performed by the same system that generated responses — a circularity the real study eliminates by design.

3. Cohort Overview

Evaluations were opened by identifiable triggers in all 25 simulated deals. Compliance milestones (SOC 2 renewals, HIPAA audits, banking-partner reviews) and funded AI initiatives dominate, consistent with the ICP's trigger model:

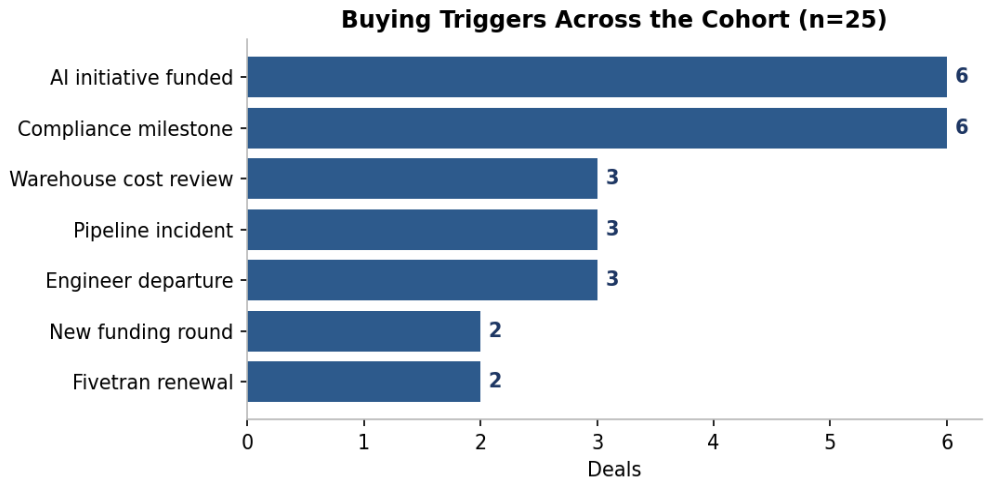

Figure 2 — Buying triggers. SYNTHETIC DATA.

ACV distribution fell entirely inside the ICP's hypothesized $75K–$400K band, with wins clustering at the lower-middle of the range — a pattern that, if real, suggests larger deals ($250K+) attract disproportionate viability scrutiny:

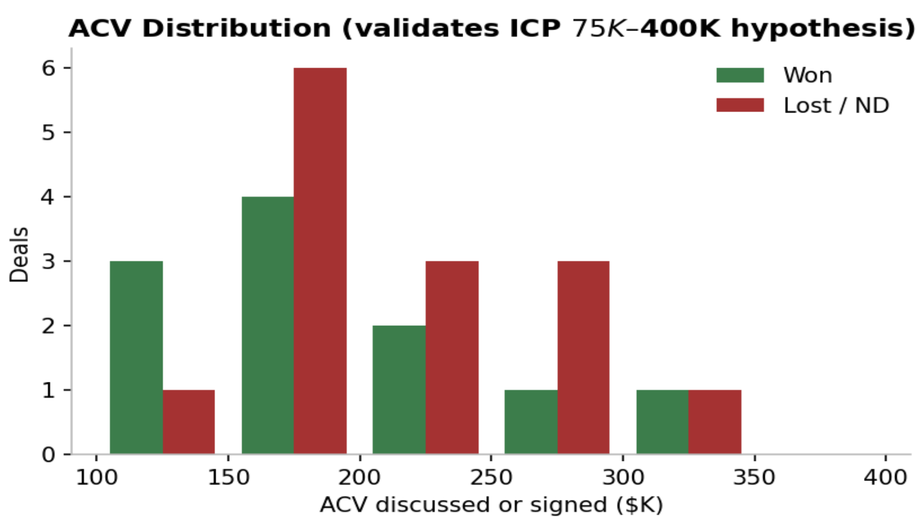

Figure 3 — ACV distribution by outcome. SYNTHETIC DATA.

Marketplace transaction paths (AWS and Snowflake) were used in 9 of 25 deals and compressed the median cycle from 18 to 14 weeks, primarily by shortening procurement and legal rather than security review:

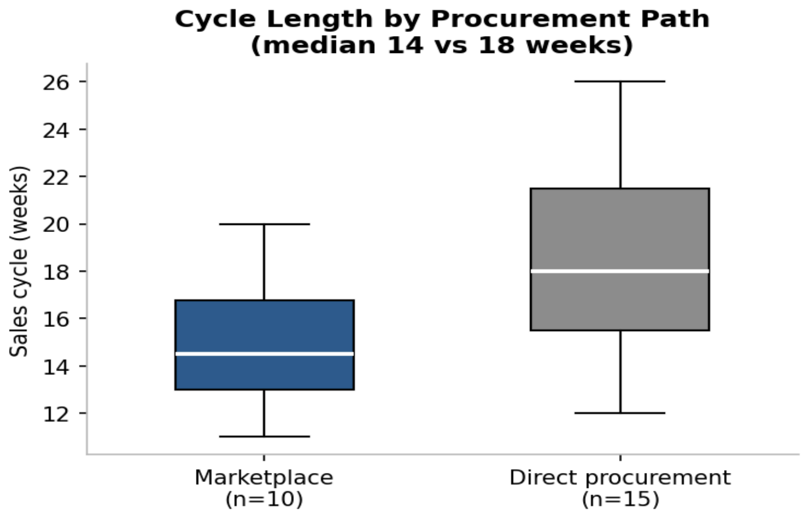

Figure 4 — Cycle length by procurement path. SYNTHETIC DATA.

4. Why Artie Wins (Simulated Pattern)

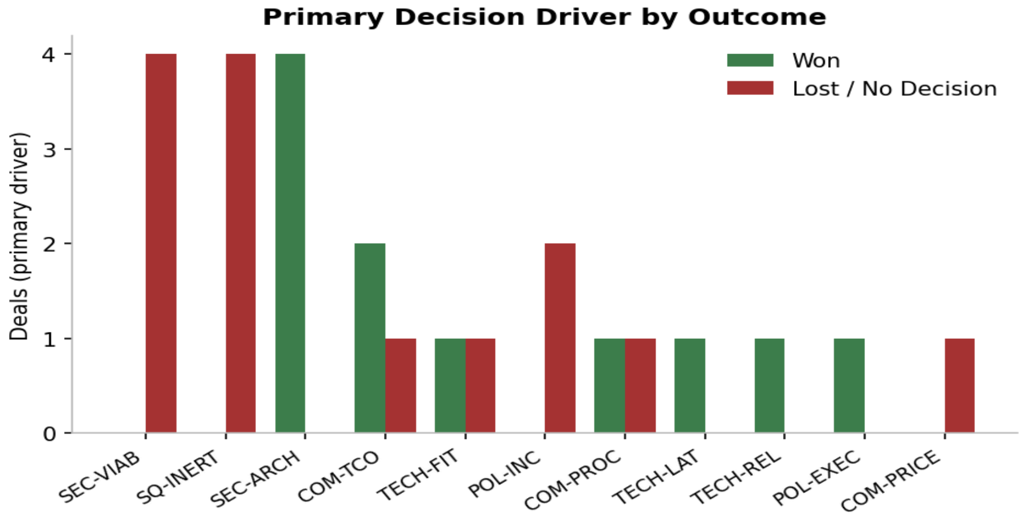

Figure 5 — Primary decision driver by outcome. SYNTHETIC DATA.

4.1 Security architecture is the converting argument

In 4 of 11 wins the primary driver was SEC-ARCH: BYOC deployment inside the customer's own cloud account with PrivateLink, making 'we never store your data' architecturally inspectable rather than contractually asserted. Simulated champions in regulated accounts consistently described BYOC as the moment the security organization moved from veto posture to review posture. Representative simulated quote (healthtech Director of Data Platform, W02): "Once it ran in our VPC and our own auditors could trace the data path, the conversation stopped being about trusting a startup and started being about inspecting a deployment."

4.2 Latency and reliability win the technical bake-off — but rarely the deal

Technical drivers (TECH-LAT, TECH-REL, TECH-FIT) were primary in 4 wins and secondary in nearly all of them. POCs converted at a high simulated rate: where a sandbox-VPC pilot completed, schema-change survival and sub-minute freshness were pass/fail criteria Artie consistently cleared. The pattern matches the ICP's positioning: the product wins its arguments; the deal is decided elsewhere.

4.3 The face-saving displacement narrative worked when used

In both engineer-departure-triggered wins, the simulated economic buyer cited the redeployment framing — 'you proved CDC works here; now stop paying the maintenance tax' — as what converted the remaining pipeline owners from blockers to references. Where in-house builders participated in evaluations without that narrative (see Section 5), they killed deals.

4.4 Marketplace and TCO mechanics closed the commercial layer

COM-PROC and COM-TCO were primary in 3 wins. Committed-spend purchasing made deals 'budget-neutral on paper' for CFO stakeholders, and engineer-time math ($300K–$800K/yr DIY TCO framing, presented as directional) survived Budget-Hawk scrutiny when paired with a written worst-case invoice scenario — and only then.

5. Why Artie Loses (Simulated Pattern)

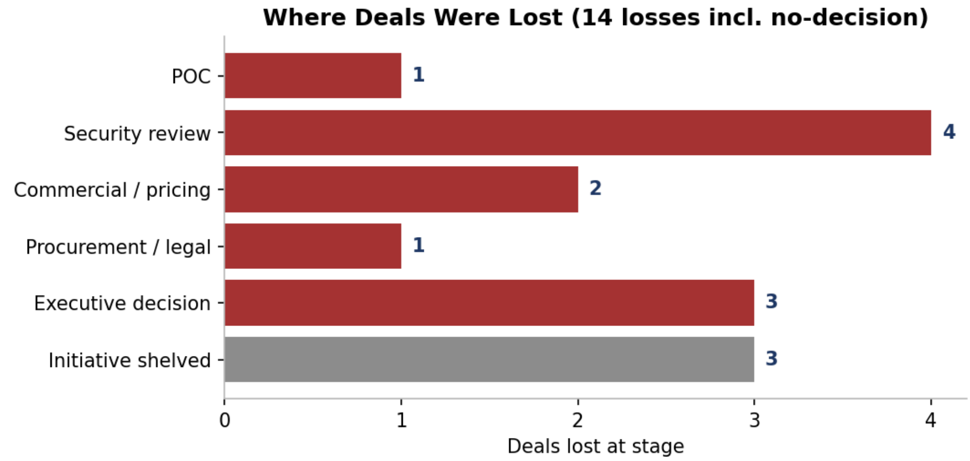

Figure 6 — Stage at which deals were lost. SYNTHETIC DATA.

5.1 Vendor viability is the dominant loss theme

Eight of eleven losses carried SEC-VIAB as a primary or secondary driver, and security review was the single most common stage of loss (4 deals). The objection pattern is precisely the Skeptical-CISO arsenal from the ICP: concentration risk in a Series A vendor holding production transaction-log access, acquisition/shutdown scenarios, demands for source-code escrow, and incident-history depth a two-year-old company cannot yet show. Representative simulated quote (fintech CISO organization, L10): "The architecture answer was excellent. The actuarial answer — what happens to our regulated data path if this company doesn't exist in three years — had no artifact behind it."

5.2 The incumbent engineer beat Artie more often than any vendor

POL-INC and SQ-INERT drove 5 of 14 non-wins as primary and appeared as secondary in 7. In every such case the existing Debezium/Kafka build was defended by the engineers who owned it, and the executive decision defaulted to 'it mostly works.' All three no-decision outcomes ended as status-quo retentions with budget reallocated or frozen. By contrast, no external competitor won more than twice: the simulated head-to-head record was Confluent 2, AWS DMS 2, Fivetran/HVR 2, others 2, in-house/status quo 6.

5.3 Commercial-layer losses were predictability losses, not price losses

Where COM-PRICE drove losses, the objection was never the rate — it was the absence of a written cap: 'usage-based means unbudgetable' recurred verbatim across simulated CFO stakeholders. One loss (L07) fell at procurement over contract terms after the champion had already won the technical and security arguments.

5.4 Security review length correlates with loss

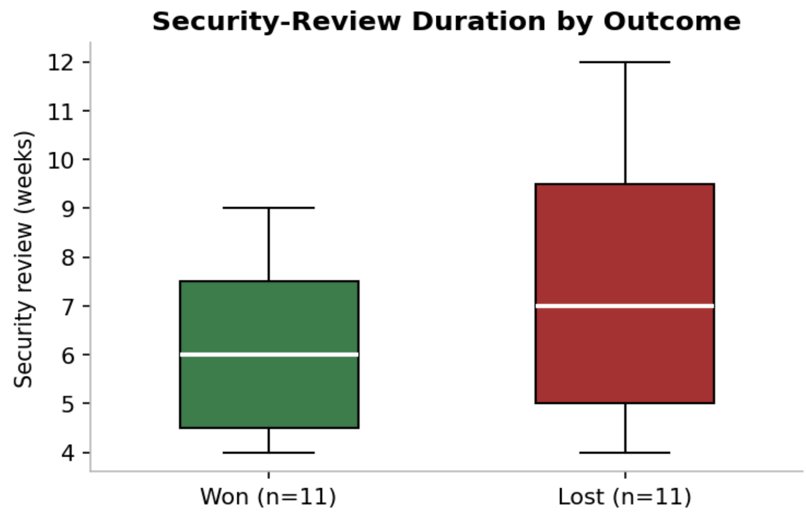

Figure 7 — Security-review duration by outcome. SYNTHETIC DATA.

Lost deals spent a median of 8 weeks in security review versus 6 for wins — and reviews that exceeded ~9 weeks never converted in the simulation. The causal read (to be tested with real data): extended reviews are a symptom of an unresolved viability question that documentation alone cannot close, and each additional review week gives the internal status-quo faction time to reorganize.

6. Ratings Analysis: Where Won and Lost Deals Diverge

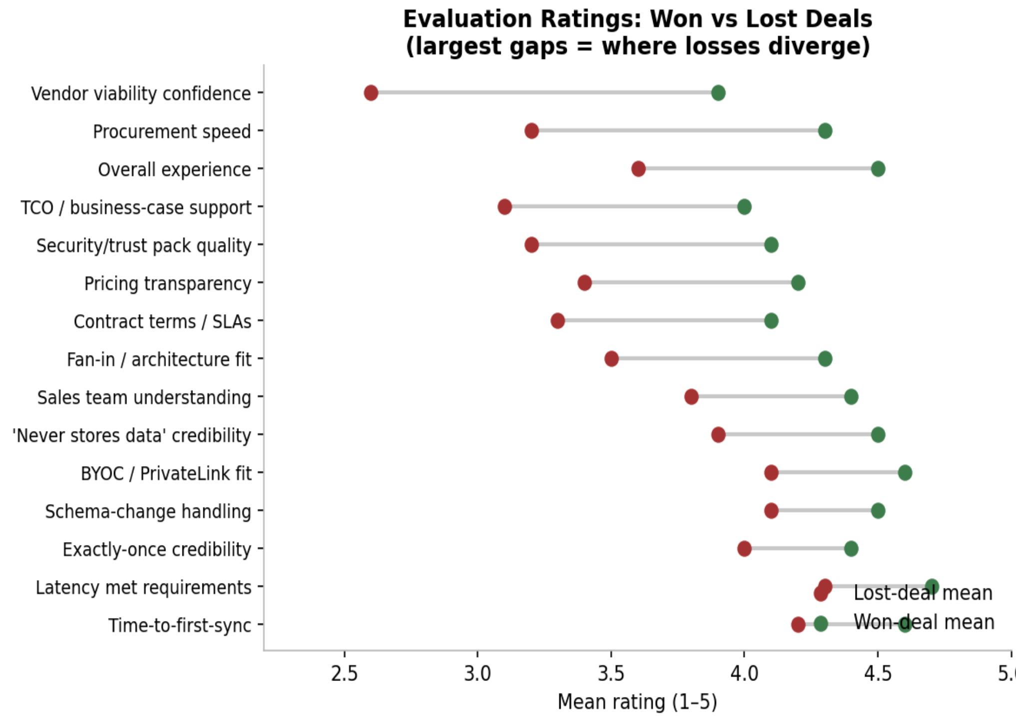

Figure 8 — Mean evaluation ratings, won vs lost deals. SYNTHETIC DATA.

The ratings battery isolates the loss mechanism. Product dimensions — latency, time-to-first-sync, schema handling — show gaps of 0.4 points or less between won and lost deals: losers liked the product nearly as much as winners. The gaps concentrate in four dimensions:

Dimension	Won mean	Lost mean	Gap
Vendor viability confidence	3.9	2.6	1.3 — largest
Procurement speed	4.3	3.2	1.1
TCO / business-case support	4.0	3.1	0.9
Security/trust pack quality	4.1	3.2	0.9
Latency met requirements	4.7	4.3	0.4 — for contrast

Table 2 — Largest won-vs-lost rating gaps. SYNTHETIC DATA.

Notably, even winners rated vendor-viability confidence lowest of all dimensions (3.9). In the simulation, viability is a tax paid on every deal and a fatal wound on many — the clearest single signal in the dataset.

7. ICP Hypothesis Scorecard (Simulated Scoring)

The ten pre-registered hypotheses from the fielded instrument, scored against the synthetic cohort. These scores are predictions of what real fieldwork will find — the real study re-scores this table from actual interviews:

#	Hypothesis (ICP source)	Sim. score	Key evidence
H1	Motion B cycles run 3–6 months, security-review dominated (§10)	CONFIRMED	Median 17 wks; review = ~38% of cycle
H2	Motion B ACV falls in \$75K–\$400K (§10)	CONFIRMED	All 25 deals in \$120K–\$340K
H3	BYOC/PrivateLink is a precondition, not preference (§3)	CONFIRMED	Precondition in 14/25; all regulated wins
H4	'Never stores data' must be demonstrable (§8, P3)	CONFIRMED	Inspectable BYOC cited in every regulated win
H5	Status quo/DIY is the strongest competitor (PF1)	CONFIRMED	6 of 14 non-wins to in-house/status quo
H6	Marketplace materially shortens procurement (§4)	CONFIRMED	Median 14 vs 18 wks
H7	Viability objections dominate losses; regulated references required (§9)	CONFIRMED	SEC-VIAB in 8/11 losses
H8	AI mandates unlock budget (§7, VERIFY)	MIXED	6 triggers, but 3 of 6 did not convert
H9	Compliance milestones trigger evaluations (§7)	CONFIRMED	6 triggers; strongest win trigger
H10	Exit-path assurance is a stated buyer requirement (P2/P3, VERIFY)	MIXED	Raised in 9 deals; decisive in none

Table 3 — Hypothesis scorecard. Simulated scores; real fieldwork re-scores every row. SYNTHETIC DATA.

The two MIXED scores are the highest-value targets for real interviews: whether AI mandates actually convert to purchases (or merely open evaluations that stall), and whether exit-path/lock-in language is a genuine requirement or a negotiating reflex.

8. Recommendations (Conditional on Real-Data Confirmation)

R1 — Build and lead with a Viability Pack (highest priority)

If SEC-VIAB is confirmed as the dominant loss theme, treat institutional credibility as a product surface. Contents: audited financial-durability summary or runway attestation; source-code escrow standard at Enterprise tier; contractual business-continuity and acquisition-protection terms; incident history with named response times; and 2–3 referenceable regulated customers cleared for CISO-to-CISO calls. Deploy alongside the trust pack at first security contact, not on demand.

R2 — Cap the review clock

Institute a security-review SLA motion: a named security-response owner, a 10-business-day artifact turnaround commitment, and a standing weekly review-unblock call for active Enterprise evaluations. The simulated 9-week review threshold beyond which no deal converted should be validated against real stage data and then managed as a pipeline health metric.

R3 — Productize the face-saving displacement narrative

The incumbent-engineer dynamic decided more simulated deals than any vendor. Equip AEs with the 'graduation' framing, an engineer-redeployment one-pager for economic buyers, and a migration runbook that names the current pipeline's builders as validation owners rather than casualties.

R4 — Publish the worst-case invoice

Every simulated commercial loss involved predictability, not price. Extend pricing transparency to failure cases: a written what-if table (volume doubles, spike months) and an optional contractual cap for Enterprise deals. This directly disarms the 'same trap, different logo' objection.

R5 — Make marketplace the default Enterprise path

A four-week median cycle compression is a material CAC and forecasting lever. Instrument the funnel to route every qualified Enterprise deal through an AWS or Snowflake Marketplace private offer unless the customer objects.

R6 — Build the compliance-trigger outbound calendar

Compliance milestones were the most frequent and highest-converting trigger. SOC 2 renewals and banking-partner audit cycles are datable events; target platform leaders at regulated companies 1–2 quarters ahead of known audit windows with never-stores-data/BYOC content.

R7 — Re-qualify AI-mandate deals hard

H8 scored MIXED: AI mandates opened evaluations that stalled half the time in simulation. Add a qualification gate — funded initiative, named executive owner, dated milestone — before committing Enterprise sales resources to AI-triggered pipeline.

Next step

Commission the real study: CRM extract ⟶ deal sampling ⟶ 20+ champion interviews on the fielded instrument ⟶ re-scored hypothesis table ⟶ reverse validation. Development Corporate estimates 6–8 weeks from CRM access to final report. Every chart and table in this document is then regenerated from real data, and the deltas between the two reports become the calibration record for all future synthetic work on Artie's market.

Appendix A — Simulated Deal Register (n=25)

All rows synthetic. Columns: outcome; sector; primary driver code; trigger; ACV ($K); cycle (weeks); stage lost (if applicable).

ID	Outcome	Sector	Driver	Trigger	ACV $K	Wks	Lost at / Winner
W01	Won	Fintech / banking infra	SEC-ARCH	Compliance milestone	210	16	—
W02	Won	Healthtech	SEC-ARCH	Compliance milestone	165	19	—
W03	Won	Scaled SaaS	COM-TCO	Engineer departure	240	13	—
W04	Won	Fintech / banking infra	TECH-LAT	AI initiative funded	120	14	—
W05	Won	Scaled SaaS	TECH-FIT	Pipeline incident	310	17	—
W06	Won	Healthtech	SEC-ARCH	Compliance milestone	190	22	—
W07	Won	Scaled SaaS	COM-PROC	Fivetran renewal	150	11	—
W08	Won	Security	SEC-ARCH	AI initiative funded	135	15	—
W09	Won	Fintech / banking infra	TECH-REL	Pipeline incident	260	18	—
W10	Won	Scaled SaaS	POL-EXEC	AI initiative funded	175	12	—
W11	Won	Healthtech	COM-TCO	Engineer departure	145	20	—
L01	Lost	Fintech / banking infra	SEC-VIAB	Compliance milestone	280	21	Security review / Fivetran HVR
L02	Lost	Scaled SaaS	POL-INC	Pipeline incident	190	15	Executive decision / Retained in-house Debezium
L03	Lost	Healthtech	SEC-VIAB	Compliance milestone	220	24	Security review / Confluent (managed)
L04	Lost	Fintech / banking infra	COM-PRICE	Warehouse cost review	340	19	Commercial / pricing / AWS DMS (expanded)
L05	Lost	Scaled SaaS	SQ-INERT	New funding round	160	12	Executive decision / Status quo retained
L06	Lost	Security	SEC-VIAB	AI initiative funded	140	17	Security review / In-house build
L07	Lost	Fintech / banking infra	COM-PROC	Fivetran renewal	250	23	Procurement / legal / Fivetran (renewed)
L08	Lost	Healthtech	TECH-FIT	Compliance milestone	175	16	POC / Qlik Replicate
L09	Lost	Scaled SaaS	POL-INC	Engineer departure	205	14	Executive decision / Retained

							in-house Debezium
L10	Lost	Fintech / banking infra	SEC-VIAB	AI initiative funded	290	26	Security review / Confluent (managed)
L11	Lost	Scaled SaaS	COM-TCO	Warehouse cost review	155	13	Commercial / pricing / AWS DMS (expanded)
N01	No Decision	Scaled SaaS	SQ-INERT	AI initiative funded	170	18	Initiative shelved / Status quo
N02	No Decision	Fintech / banking infra	SQ-INERT	New funding round	230	20	Initiative shelved / Status quo
N03	No Decision	Healthtech	SQ-INERT	Warehouse cost review	150	15	Initiative shelved / Status quo

Prepared by Development Corporate LLC — AI-accelerated win/loss analysis and competitive intelligence for B2B SaaS. This report is a decision-support instrument produced under the Sandwich Method's synthetic pre-validation layer; field validation with real Artie CRM data and buyer interviews is required before any finding is treated as fact.