VEGA STUDIO
RESEARCH
PAPER 02
Q4 2026

The Messaging Capital Leak

Why most go-to-market waste is decided before the first sales call, and how to price it

94% of buying groups have already ranked a preferred vendor before first contact with any seller	**77%** and they buy from that preliminary favorite 77% of the time

6sense, B2B Buyer Experience Report (2025)

EXECUTIVE SUMMARY

Most companies that are burning go-to-market capital believe they have an execution problem.

They usually have an arithmetic one.

Somewhere upstream of the sales call, the company decided which accounts deserve its attention. That decision was almost always made using a variable that is convenient to target: headcount, industry, job title, funding stage. And in a surprising number of companies, that variable appears nowhere in the formula by which the company actually makes money.

Once that gap exists, every downstream investment inherits it. More sellers, more content, and more budget do not close it. They multiply it.

Why nobody catches it. Your dashboard counts deals. Your revenue line counts revenue. Neither counts which deals. So a customer who repays your acquisition system in three weeks and a customer who will never repay it appear on the same chart, as the same event, in the same color.

WHAT IT COSTS

In the audit behind this paper, 62% of a company's acquisition budget was going to a segment that returned under 2.5% of its results, at 67 times the cost per click. That was the visible half. The invisible half was worse: the distance between its best and worst customers, measured in months to repay acquisition, was greater than two hundred to one.

The diagnostic. Four symptoms, each testable with one question and priceable with one calculation. We call the set the ICP Drift Audit. It runs in an afternoon on data you already have.

The measurement. One number, which we call the Payback Spread: the distance between the customers who repay your acquisition system fastest and the ones who may never repay it at all. It is very likely the most useful number about your business that nobody currently owns.

The intervention takes about three weeks, because the analysis is fast and the decision is not.

The complete framework, the four tests, the calculations, and the three-week method are all in this paper. There is no withheld step and nothing gated. If you can run it yourself, run it yourself.

The system, in one view

Problem, diagnostic, symptoms, measurement, intervention

THE PROBLEM	Messaging Capital Leak
THE DIAGNOSTIC	ICP Drift Audit
THE FOUR SYMPTOMS	Proxy Drift; Discount Reflex; Concentrated Churn; Unworked Ceiling
THE MEASUREMENT	Payback Spread
THE INTERVENTION	Three weeks: diagnosis, architecture, roadmap

SECTION 01

What the leak costs

The arithmetic nobody runs

Every company can tell you its customer acquisition cost.

Almost none can tell you its acquisition cost by segment, which is the only version of the number that means anything.

Consider a platform business we audited in 2026. It sells through partner distribution and earns a recurring commission on the volume its partners generate. Its revenue formula was simple and never in dispute:

$$\text{Revenue} = \text{commission rate} \times \text{partner user volume}$$

Anyone in the company could have recited it.

Its paid acquisition, meanwhile, was targeting by industry and by company headcount, in a band of 50 to 5,000 employees.

Headcount does not appear in that formula. It is not a weak proxy for user volume. It is not a proxy at all. A financial technology company with 1.5 million users may run on 180 employees. A traditional institution with 40,000 employees may have fewer digitally active users than that company does.

The targeting variable and the revenue variable were unrelated, and had been for over a year.

FIGURE 1

Nearly two thirds of the budget produced under one fortieth of the output

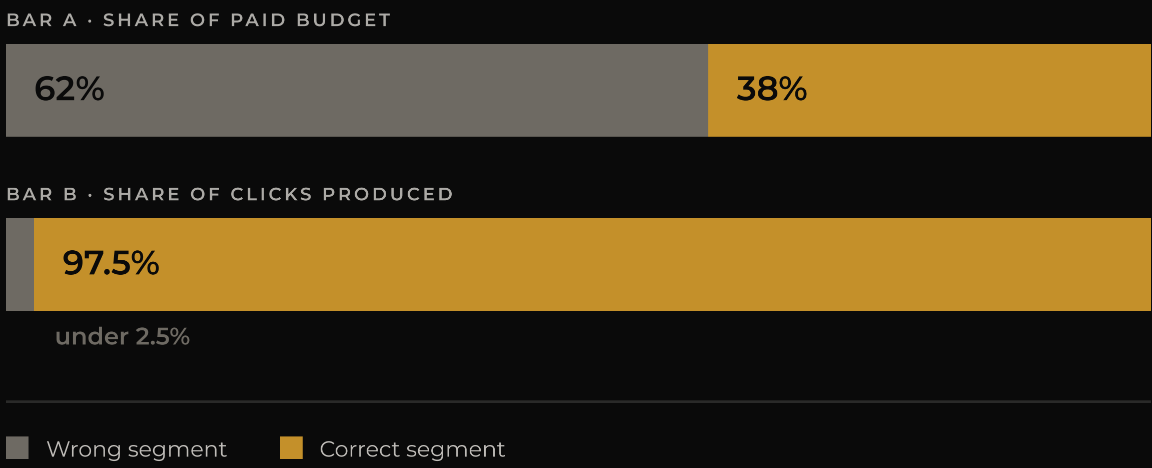

Own audit, anonymized, May 2026. Single month, single account.

FIGURE 2

The same money, pointed correctly, was 67 times cheaper

84×

CLICK-THROUGH RATE

5.02% correct segment vs 0.06% wrong segment

67×

COST PER CLICK

US$0.020 correct segment vs US$1.343 wrong segment

Own audit, anonymized, May 2026.

Note what this is not.

It is not a creative problem, a channel problem, or a bidding problem. The same creative, in the same channel, on the same budget, performed extraordinarily well against the correct audience in the same month.

The only broken variable was who the money was pointed at.

The year before, the same account had spent for six months and recorded zero conversions for the entire year, because nothing downstream of the click was instrumented. Money leaving, nothing measured, nobody able to say whether it worked.

The Payback Spread

The budget figures above are the visible half of the leak. This is the half that does not appear on any dashboard, and it is the one to take into your next board meeting.

Payback Spread measures the distance between the customers who repay your acquisition system fastest and the ones who may never repay it at all.

To calculate it:

> Payback Spread = months-to-payback of your worst-fit account ÷ months-to-payback of your best-fit account

You need two inputs you already have: acquisition cost and account-level revenue. Nothing else.

FIGURE 3

Two customers, same team, same quarter, 200 times apart

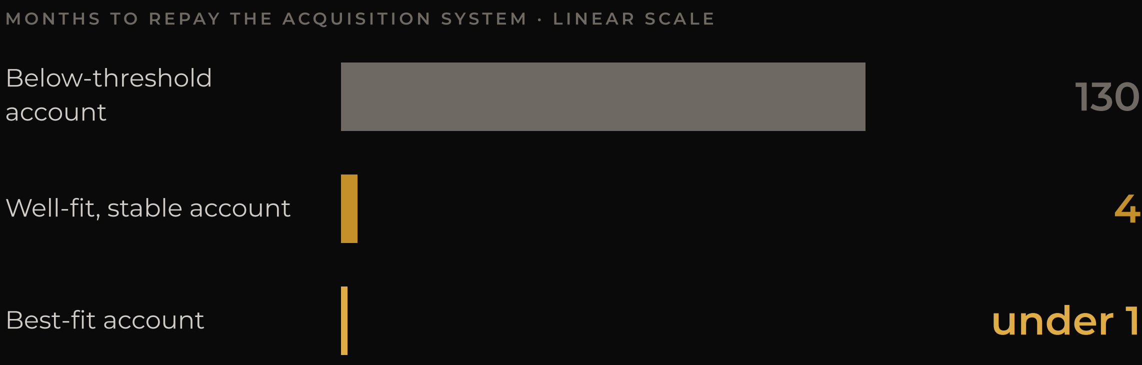

SPREAD BETWEEN THE EXTREMES: MORE THAN 200 TO 1

Own audit, anonymized, 2026. Ratios derived from client-confirmed account economics.

Sit with that.

Two customers, sold by the same team, using the same deck, closed in the same quarter. One repays the entire cost of the acquisition system in under a month. The other takes almost eleven years, which in practice means never, because the account will churn, the product will change, or the company will have pivoted long before.

Both of them show up in the pipeline as a win.

That is why conventional reporting cannot see this. Volume metrics are indifferent to composition. A quarter of thirty low-fit wins and a quarter of thirty high-fit wins are the same quarter on the chart and different companies in reality.

Why the timing makes it worse

Two external findings explain why this failure costs more than it used to.

The first is when the decision gets made. According to 6sense's B2B Buyer Experience Report (2025), 94% of buying groups have already ranked a preferred vendor before first contact with any seller, and they buy from that preliminary favorite 77% of the time.

> If you are positioned for the wrong buyer, you are not late in the process. You were never in it.

The second is duration. ICONIQ's State of Go-to-Market 2026 puts the typical B2B sales cycle at roughly 19 weeks, extending to about 24 weeks for deals above US$100K. A misqualified enterprise opportunity consumes close to six months of a seller's finite capacity before resolving into nothing.

That is not a wasted call. It is a wasted half-year of a person you are paying.

SECTION 02

Why the obvious fixes make it worse

Three responses are close at hand when pipeline underperforms. All three are rational. All three scale the leak.

Hire more sellers.

A sales organization is a multiplier applied to a qualification rule. Doubling it doubles the rate at which capital converts into low-fit revenue. The pipeline chart improves. The payback distribution gets worse. This is the most expensive of the three because it is the hardest to reverse.

Buy more traffic.

In the audit above, demand generation was performing at a 7.86% click-through rate on its best format, at US$0.025 per click. Strong by any standard. The demand was real. It simply was not attached to anything that could become revenue.

Add tooling.

Attribution platforms, intent data, and scoring models answer "which of these leads is most engaged." None answers "should this account have been in the set at all." A scoring model built on the wrong qualifier ranks the wrong accounts more precisely.

Each treats insufficient output as the constraint, when the constraint is definitional and sits upstream of all three.

What companies actually experience is not a bad quarter. It is several consecutive quarters of visibly harder work with flat compounding, which is harder to diagnose and considerably more demoralizing.

The reason it survives is not incompetence

In the audit above, the executive sponsor asked two questions inside his own company. What is a new partner actually worth to us. How many qualified leads did last year's advertising produce.

He got no usable answer to either. He eventually assembled the numbers himself, by hand.

Those questions had no owner.

That is the ordinary condition of a growing company, not a failing one. The leak survives because the information needed to see it sits in three different functions and belongs to none of them.

MARKETING	FINANCE	SALES	NOBODY
holds the spend	holds the revenue	holds the account context	holds the ratio

SECTION 03

The ICP Drift Audit

Four symptoms. Each has one diagnostic question and one calculation. All four run in an afternoon on data you already have.

FIGURE 4

The four tests, and what each one is worth finding

SYMPTOM	THE ONE QUESTION	THE CALCULATION
Proxy Drift	Does your primary targeting variable appear in your revenue formula?	Budget share to the proxy segment × performance gap
Discount Reflex	Is the average discount on your last ten deals higher than on the ten before?	Average discount × gross margin × deal volume
Concentrated Churn	Do your last ten churned customers resemble each other more than the ten who stayed?	Lost recurring revenue + acquisition cost + service cost
Unworked Ceiling	When did you last open an opportunity worth 3× your average, and who owned it?	Account value × conservative win rate × periods unattempted

Symptom 1 · Proxy Drift

You are targeting, scoring, or prioritizing on a variable that does not appear in your revenue formula. Usually a convenient variable rather than a true one, because convenient variables are the ones every targeting platform offers.

> **Test:** *Write your revenue formula in one line. Does your primary targeting variable appear in it?*

Cost: share of acquisition budget going to the proxy segment, multiplied by the performance gap between that segment and your true one.

THE UNCOMFORTABLE COROLLARY

Sometimes the true qualifier cannot be targeted at all. In the audit, the real qualifier was partner user base above one million, and no advertising platform lets you segment on that. The honest conclusion was that broad targeting could not serve this business at any budget, and acquisition had to be rebuilt around a named account list assembled by hand.

That is a harder answer than "optimize the campaign." It was the correct one.

Symptom 2 · The Discount Reflex

Discount has quietly become the closing tool. Not a concession at the end, but the thing that actually moves deals from late stage to closed.

> **Test:** *What was the average discount on your last ten closed deals, and on the ten before those?*

A rising number is often evidence that positioning is weakening and price is doing the work positioning stopped doing.

Cost: average discount percentage × gross margin × deal volume, annualized. For most B2B companies this is the single largest number in the audit, and the one most likely to be dismissed as "just how the market is."

Why it is a positioning symptom rather than a pricing one. A buyer who has already ranked you first does not need to be discounted into the deal. Discount is frequently what you pay to overcome a position you did not earn earlier.

Symptom 3 · Concentrated Churn

Your churn is not evenly distributed. It is concentrated in a segment, and that segment is one you decided to sell to.

> **Test:** *List the last ten customers who left. Do they resemble each other more than they resemble the ten who stayed?*

If they do, you did not lose those customers. You acquired them.

Cost: recurring revenue lost from the concentrated segment over twelve months, plus the acquisition cost spent to win them, plus the service cost to keep them while they were leaving. Most companies count only the first of the three, which understates the number by roughly half.

Symptom 4 · The Unworked Ceiling

There is a class of account worth several times your average, you are technically capable of serving it, and nobody has opened one in a long time. Not because it was tried and lost. Because it was never attempted.

> **Test:** *When did your team last open a new opportunity worth three times your average deal, and who owned it?*

A vague answer here usually indicates a positioning gap rather than a capability gap. The company does not believe it can credibly show up in that conversation, so it never has the conversation.

Cost: this is opportunity cost, so it belongs in the audit as a range rather than a point. Value of one such account × a conservative win rate × the periods it went unattempted. In the audit above, a single best-fit account was worth roughly eighteen times the annual cost of the entire acquisition system.

Scoring it

FIGURE 5

The audit scorecard

SYMPTOM	PRESENT?	ESTIMATED ANNUAL COST
Proxy Drift		
Discount Reflex		
Concentrated Churn		
Unworked Ceiling		
Total leak		
Total as a percentage of annual revenue		

Two rules for filling this in honestly.

Any figure you cannot source gets marked not measured rather than estimated. An audit with three real numbers and one blank is more useful than one with four plausible numbers.

And run the total against annual revenue as a percentage. That percentage is the sentence you take to the board, because it is the only version a finance function will act on.

SECTION 04

The three-week method

FIGURE 6

Three weeks, because the analysis is fast and the decision is not

	OBJECTIVE	DELIVERABLE
Week 1 Diagnosis	Establish what is true, separately from what the company believes	Positioning audit, priced
Week 2 Architecture	Make the decision, then make it durable	Messaging architecture and one-page decision governance
Week 3 Roadmap	Sequence the change so it survives	90-day roadmap, quality metric live, one-page playbook

Week 1 · Diagnosis

Write the revenue formula in one line. Get agreement from finance and from sales independently. If the two functions write different formulas, stop. That disagreement is the whole project and everything else is downstream of it.

Census the accounts. Every customer from the last twenty-four months, with acquisition cost, revenue to date, tenure, and churn status. Compute months-to-payback for each. Sort descending. Your Payback Spread falls out of this table.

Run the four-symptom audit.

Interview separately. Founder or chief executive, whoever owns revenue, and at least two sellers. Ask each independently to describe the ideal customer. Write the answers verbatim and do not reconcile them yet. The divergence is your evidence.

Count the definitions in circulation. Website, sales deck, last ten proposals, last quarter of outbound. You are looking for one thing: how many distinct definitions of the customer are currently live.

Week 2 · Architecture

Define the qualifier. One variable, drawn from the revenue formula, that determines whether an account is in or out. It must be observable from outside the company before a sales conversation, or it is not usable. This is the hardest hour of the three weeks and everything depends on it.

Build the messaging architecture. One document, one source of truth: category, audience, value proposition, proof structure, language standards. It exists so a new employee, an agency, or a freelancer can make a message decision without escalating.

Design decision governance. One page. Who owns which message decision, and what does and does not go up to leadership. This is the most frequently skipped step in the method. Skipping it is why most positioning work does not survive the following quarter.

Week 3 · Roadmap and handover

Build the ninety-day roadmap. Sequenced, with named owners, ready to execute on day one. Not a Gantt chart. An operating document ordered by what unblocks what.

Re-target acquisition against the new qualifier. In the audit above this meant abandoning platform targeting entirely for a hand-built named-account list. It belongs in week 3 precisely because it is the first thing that gets quietly reversed under pressure if it is not written down with an owner.

Define the quality metric. Not lead volume. The percentage of new opportunities that meet the qualifier. On the dashboard, next to volume, with the same prominence.

Hand over, then compress everything into a one-page playbook.

Why this is hard to run on yourself

We have given away the entire method, which raises a fair question.

The answer is not that the analysis is difficult. It is not. Any competent operator can build the payback table and score the four symptoms in a day.

The framework is simple. Organizational decision-making is not.

The difficulty shows up in three specific places, and none of them are analytical.

> **The revenue formula usually has two versions and no referee.** Finance writes one. Sales writes another. Both are defensible, both are partly right, and neither function can adjudicate the other's, because each has an interest in the answer. Somebody with no stake in either version has to sit in the room and force the reconciliation. This is where most self-run audits stall, and they stall quietly: the exercise gets rescheduled rather than refused.

> **The person best placed to run the audit is often the person whose earlier decision created the drift.** The targeting was set by someone. The segment was chosen by someone. Asking that person to produce the evidence that the choice was wrong is asking a great deal. It is not a character flaw that it rarely happens. It is the predictable result of who holds the pen.

> **The governance page requires someone to lose authority.** Writing down who owns which message decision means somebody stops owning decisions they own today. That is a political act, not an analytical one, and internal actors pay a real cost for proposing it. Outsiders do not.

Information is cheap. Judgment under organizational pressure is not.

If you have a leadership team that can hold a hard conversation about a formula and reassign decision rights without it becoming personal, run this yourself. Many companies can. The ones that cannot are not less capable. They are subject to the ordinary physics of an organization where everyone in the room has a history with the answer.

One warning about measurement

AGREE THIS BEFORE YOU START

If your sales cycle is long, do not measure this work by closed deals in the first two quarters. A complex partnership or enterprise deal commonly closes in six to eighteen months, so most of a first-year pipeline closes in year two.

Agree before you start that year one is measured by qualified opportunities opened and pipeline in conversation, not by revenue closed.

This is not a soft target. Getting it agreed in writing beforehand is what prevents the program being cancelled in month six, immediately before it matures. We have seen that cancellation happen, and it is the most avoidable failure in this entire method.

SECTION 05

Two engagements

We publish outcomes that stayed with the company rather than campaign metrics. Client identities are withheld unless we hold written permission to name them. Where a figure was never measured, we say so rather than estimating it.

Two engagements in depth rather than four in summary, because a case without a mechanism is a testimonial, and testimonials are not evidence.

ENGAGEMENT A

A platform business selling through partner distribution

The situation. The company earns a recurring commission on the transaction volume its distribution partners generate. Revenue is a direct function of how many end users a partner has. It operates across six countries with heavily concentrated revenue: one large partner accounts for roughly a quarter of the annual total, while the smallest contribute a rounding error.

What it believed the problem was. Marketing was not producing pipeline. The internal reading was that advertising was underperforming, and the response, in the quarter before we arrived, had been to cut the media budget by half.

What the audit found. The advertising was one of the healthiest parts of the business, delivering a 7.86% click-through rate on its strongest format at US$0.025 per click.

The problem was that acquisition targeted industry plus headcount while revenue depended entirely on partner end-user volume. 62% of media budget went to a segment returning under 2.5% of clicks, at 67 times the cost per click. Underneath that sat a Payback Spread above 200 to 1, and nothing in the reporting distinguished between the two ends of it.

Why nobody had caught it. The executive sponsor had asked internally what a new partner was worth and how many qualified leads the previous year of advertising had produced. He received no usable answer to either and built the numbers himself, by hand, over messaging. The questions were not hard. They had no owner.

What changed. The qualifier was redefined around partner user base rather than headcount. Because no advertising platform can segment on user base, broad targeting was abandoned for a hand-assembled named-account list, which is a materially different acquisition motion and a harder one to staff. And one quality metric went onto the dashboard beside the volume metrics: the percentage of new opportunities meeting the user-volume threshold.

WHAT WE ARE NOT CLAIMING

This engagement is in progress. We are reporting a diagnosis and a set of decisions, not an outcome, because the outcome is not in yet. Deals in this category close in six to eighteen months, so the first honest read is a year away. We will publish it when we have it.

ENGAGEMENT B

A mobility platform

The situation. A consumer application in shared mobility, live in market, burning capital, with founders personally involved in making individual transactions complete.

What it believed the problem was. Marketing. The company was preparing to spend more on acquisition on the assumption that the product worked and awareness was the gap.

What the audit found. The founders did not have a shared account of their own business model. Two incompatible versions were operating at once: a consumer transaction business earning a small percentage of a small transaction, and an institutional business selling closed, verified communities to organizations.

Those two businesses have different customers, different price points, different products, and different definitions of success. Nobody had chosen. So the pricing logic was built for one and the product roadmap for the other, and marketing was asked to generate demand without being told which business it was generating demand for.

No amount of media spend resolves that, because the ambiguity is upstream of the spend.

What changed. Value proposition, charging model, and product scope were redefined together rather than separately, which is the only way they can be redefined, since each constrains the other two.

The outcome, and why we report it this way. Transactions began completing end to end with no founder involvement at all. Before that point the business could not operate without its owners in the loop on individual transactions.

We report that rather than a volume figure deliberately. The volumes at that stage were small, and quoting them would be flattering to nobody and informative to no one. What changed was not a number, it was a capability: the company stopped being a service its founders performed and started being a system that ran. That capability stayed after we left.

What is missing from both. Cycle length and average contract value, before and after. Those were not instrumented at the time. We are capturing them prospectively from this quarter forward and they will appear in a future edition.

SECTION 06

Three objections

"We already have our positioning defined."

Almost every company we audit does, and the document is usually good.

The question is not whether it exists. It is whether the same positioning debate has returned to the leadership table in the last two quarters. If it has, you have a document, not a system.

A document persuades. A system decides. The difference is decision governance: a written answer to who owns which message decision. Positioning work that omits governance tends to erode within about two quarters, and the erosion looks like drift rather than failure, which is why it goes unnoticed.

"Our problem is demand generation, not messaging."

Sometimes true. Here is how to tell in ten minutes.

Look at your click-through rate and cost per click against your best-fit segment specifically, not blended. If those are strong and pipeline is still weak, adding demand will make the leak larger rather than smaller.

In the audit behind this paper, demand generation was performing well above benchmark at the exact moment the account was recording effectively no qualified pipeline. Blended metrics hid it completely, which is the ordinary case rather than the exception.

"Narrowing the ICP means turning away business we close today."

This is the real objection and it deserves a real answer rather than reassurance.

Yes. It does mean that, and it should be decided with numbers in front of you rather than in principle. Run the payback table and find where months-to-payback crosses your planning horizon. Accounts on the far side of that line are not revenue. They are financed liabilities: they consume acquisition cost, service capacity, and product attention, and they return the investment on a timescale longer than your ability to plan for it.

Two honest qualifications. Below-threshold accounts sometimes carry strategic value that is real and does not appear in the payback math: a reference logo, a distribution relationship, a product signal. Take those deliberately, priced as the investments they are, rather than by default.

And narrowing does not require firing existing customers. It changes where new capital is pointed, which is a decision about the next dollar rather than the last one.

SECTION 07

Run the audit

Four questions. If you can answer all four with a number, you do not need help. If two or more are blank, the blanks are the finding.

FIGURE 7

The four questions

1	Write your revenue formula in one line. Does your primary targeting variable appear in it?
2	What was the average discount on your last ten closed deals, and on the ten before those? Is it rising?
3	Of your last ten churned customers, do they resemble each other more than the ten who stayed?
4	When did your team last open an opportunity worth three times your average, and who owned it?

Then compute one number.

Take your customer list from the last twenty-four months. Calculate months-to-payback for each account. Divide the worst by the best.

That ratio is your Payback Spread.

THE NEXT STEP

The ICP Drift Diagnostic

A structured twenty-minute working session. You bring whatever numbers you have. We calculate an initial Payback Spread with you and score the four symptoms against your business.

In twenty minutes you establish whether your constraint is downstream execution or upstream qualification. Those two problems look identical on a dashboard and have almost nothing in common as work.

You leave with the calculation either way.

fran@vega-studio.com

About this paper

Vega Studio publishes one flagship paper per quarter, from audits and real engagements rather than from opinion. Figures come from official provider documentation, public reports, or our own anonymized audits. Where something was not measured, we publish the gap rather than an estimate.

This is paper 02. Paper 01, El Huevo y la Gallina, is about email deliverability and shares this paper's underlying thesis: the problem is never where you notice it.

SOURCES

6sense, B2B Buyer Experience Report (2025)

Source of the finding that 94% of buying groups rank a preferred vendor before first contact and purchase from that favorite 77% of the time.

ICONIQ, State of Go-to-Market 2026

Source of the approximately 19-week typical B2B sales cycle, extending to roughly 24 weeks for deals above US$100K.

Vega Studio own audits, anonymized, 2025 and 2026

Source of all segment performance figures, the payback distribution, and the Payback Spread ratio.

Figures from client engagements are published as ratios or as directional outcomes. Absolute revenue figures, client identities, and third-party names are withheld.

VEGA
STUDIO

RESEARCH · PAPER 02 · Q4 2026