AMBASSADOR

WHITE PAPER

The Customer Lifecycle Operating System, orchestrated by HiroAI

The Advocacy Plateau.

And what comes after it.

Every advocacy program produces a fast lift, then flattens. The plateau is not the program failing. It is the program finishing the only audience it was built to reach. This paper sets out why the curve flattens, what has to be true architecturally for it to keep climbing, and the economics of the inflection where acquisition cost falls while revenue keeps rising.

PUBLISHED BY
Ambassador

TOPIC
Customer lifecycle maturity

AUDIENCE
CMO, CRO, CFO

getambassador.com

EXECUTIVE SUMMARY

The plateau is arithmetic, not failure.

Advocacy is the highest-return first move in customer growth, and the evidence is not anecdotal. In the most rigorous study of the question, referred customers at a European retail bank showed lifetime value at least 16 percent higher than comparable non-referred customers, contribution margins roughly 25 percent higher in the first year, and churn speed about 18 percent lower.[1] Which is why nearly every company that launches a referral program sees a genuine lift inside two quarters.

It then flattens. The near-universal response is to question the program: refresh the incentive, increase promotion, change vendors. None of it works for long, because the plateau is not a program problem. Whatever audience the program was pointed at is finite, and once everyone who will participate has participated, a saturated curve does what saturated curves do.

The revenue did not disappear. It moved to audiences nobody is asking: the customer base, the product surface, the partner channel, and the lifecycle signals that pass by unused. We call that leakage, and it is the most expensive line item in customer growth precisely because nothing about it registers as a loss.

> **THE CORE THESIS**
>
> Advocacy is the entry point, not the ceiling. Reaching the next audience is not a feature problem solved by buying a second tool. It is an architecture problem, because the second audience only produces compounding revenue when every signal about the same customer lives on one data model, with one orchestrator deciding what fires and when.

The timing matters. Paid acquisition has grown materially more expensive, search behaviour is moving to AI answers that do not pass traffic through, and platform-level signal loss has made targeting and measurement harder at the same time. Each of those trends transfers growth weight from audiences a company can buy to audiences it already owns. This paper sets out the three structural failures behind the plateau, introduces the Customer Growth Maturity Model as a diagnostic, presents the external evidence for why the shift is now rather than later, and describes a staged path that does not require buying a lifecycle on day one.

$2.4B+
Revenue attributed through the platform

25%+
Advocacy rate

3:1 to 5:1+
Return on investment

99.9%
Platform uptime

Ambassador platform figures. External research is cited throughout and listed in full in the references section.

1 Why the first program works

It is worth being precise about why advocacy produces a fast lift, because the reasons are structural and they are also the reasons the lift is finite.

1.1 The evidence on referred customers

The foundational study is Schmitt, Skiera and Van den Bulte, published in the Journal of Marketing in 2011. The authors tracked roughly 10,000 accounts at a German retail bank over 33 months and compared referred customers against non-referred customers matched on demographics and acquisition timing.[1]

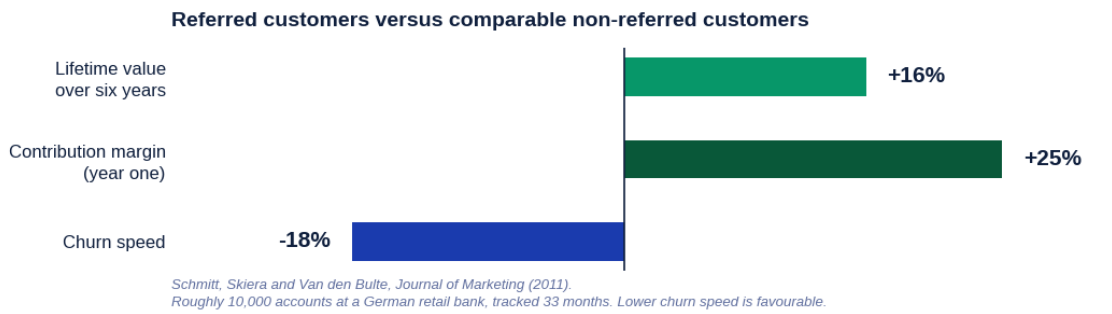

Figure 1. Referred customers outperform on value, margin and retention. The margin advantage decays over time. The retention advantage does not.

Two details in that paper deserve more attention than they usually get. First, the contribution margin advantage is largest in year one and erodes thereafter, while the retention advantage persists. Referral is therefore a durable retention lever, not only an acquisition discount. Second, the effect varies widely by segment, from roughly minus three percent to plus 36 percent, which is why the authors recommend a selective rather than a blanket referral approach.[1]

On the demand side, Nielsen's global trust research found that 88 percent of respondents trust recommendations from people they know more than any other channel, the highest-scoring source in the study.[2]

1.2 Three durable advantages

→ **Cost.** Referred acquisition is cheaper because the introduction is made by someone already inside the base rather than purchased at auction. In the bank study, acquisition cost for referred customers ran roughly 20 euros lower per customer.[1]

→ **Quality.** A referred prospect arrives with a recommendation attached, and closes and retains better than a comparable paid-sourced prospect.[1]

→ **Direction of scale.** Almost every acquisition channel gets more expensive as it gets bigger. Advocacy gets cheaper, because the producing population is one the business already pays to serve.

So the first program works, and it works for reasons that have nothing to do with which vendor was selected. What the spike proves is that a specific audience will refer you when asked. That is a valuable thing to have proven. It is also worth exactly one audience.

> *The spike is not the achievement. It is the evidence that the channel is real.*

2 Three failures, not one

At the plateau, the same three problems appear almost every time. None of them is a program problem, which is why adjusting the incentive never resolves them.

2.1 The audience runs out

This one is arithmetic. If the program asked employees, there is a fixed number of employees and a fixed share who will ever participate. If it asked the best customers, the same holds. Saturation is the expected outcome of a successful program pointed at a bounded population.

WHERE THE REVENUE ACTUALLY GOES

Your first program captures one audience. The rest is still there.

CAPTURED	NOT ASKED	NOT ASKED	NOT ASKED	NOT ASKED
The audience you asked	**Your customers and clients**	**Inside your product**	**Partners and resellers**	**Lifecycle signals NPS, reviews, renewal**

TOTAL REFERRAL REVENUE AVAILABLE TO YOU

The plateau is not the program failing. It is the program finishing the only audience it was built to reach.

That is leakage. The revenue did not disappear. It moved somewhere nobody is asking for it.

Figure 2. The first program captures one audience. The rest of the available referral revenue is still there, unasked.

2.2 The data stops at the tool

A referral tool captures a referral and hands off a lead. The trail ends there. Advocacy sits in one silo, retention in another, surveys and reviews in a third, paid media in a fourth. Nothing shares a customer record.

This is not a niche complaint. Gartner's cross-functional customer data research found that only 14 percent of organisations had achieved a single view of the customer.[3] Survey work by the CDP Institute across 200 brands found 68 percent still struggling with siloed data preventing a unified view.[4]

The operational consequence is that reward eligibility gets rebuilt by hand out of CRM reports every cycle, and attribution becomes a spreadsheet exercise. The strategic consequence is worse: even when the program is working, nobody can demonstrate what it returned without assembling the answer first. A number that has to be assembled is a number that loses budget arguments.

2.3 Each next move costs another vendor

A rewards storefront is a vendor. A leaderboard and participant dashboard is often a second. Surveys and NPS, a third. A customer-facing program alongside an internal one, frequently a fourth contract with a fourth integration. Each lands in the same program owner's budget, and that line grows every year the program succeeds.

The category has been built to encourage exactly this. The marketing technology landscape reached 14,106 products in 2024, up 27.8 percent on the prior year.[5] Meanwhile Gartner's marketing technology survey found utilisation of stack capability had fallen to 33 percent on average in 2023, from 42 percent in 2022 and 58 percent in 2020, with overlap between solutions named as a driver. Utilisation recovered to 49 percent by 2025, but only 15 percent of organisations qualified as high performers.[6]

THE ROOT CAUSE

The problem is not that point solutions are poor software. Several are excellent. The problem is that both the pricing unit and the data model are scoped to a single program. Scaling the outcome scales the bill and never scales the intelligence. The pricing unit is the trap.

3 Why the shift is now

The argument in this paper would have been true five years ago and would have been easy to defer. It is harder to defer now, because three independent trends are all moving in the same direction: away from growth a company can buy and towards growth it has to earn from the customers it already has.

3.1 Paid acquisition is more expensive and slower to pay back

Benchmarkit's 2025 study of roughly 1,000 B2B software companies found the median new customer acquisition cost ratio rose 14 percent during 2024, to two dollars of sales and marketing spend per dollar of new annual recurring revenue. Fourth-quartile companies spend $2.82. Top-quartile companies are near one dollar.[7] Median payback on that spend now runs around 16 months, with the bottom quartile at 24 months or worse.[8]

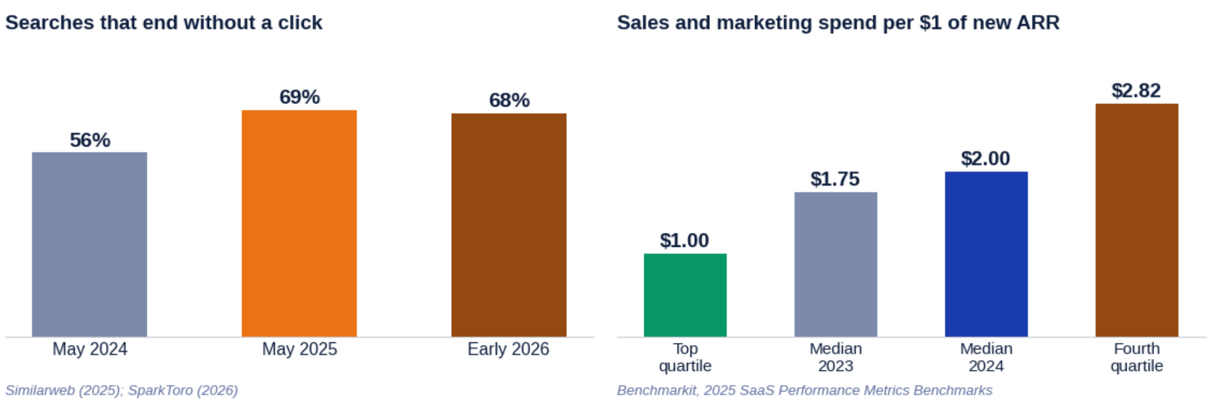

Figure 3. Two independent pressures. Search increasingly resolves without a click, while the cost of buying a dollar of new revenue continues to rise.

Platform pricing reinforces it. Meta reported average price per ad up 6 percent year over year in the fourth quarter of 2025 and 9 percent for the full year, on rising impression volume.[9] An analysis of more than 16,000 Google Ads campaigns found cost per click increased for 87 percent of industries year over year.[10] Commercial estimates put overall acquisition cost inflation at roughly 60 percent over five years, though that figure is vendor-produced rather than independently audited.[11]

3.2 Search is resolving without a click

The most methodologically careful work here is from the Pew Research Center, which analysed 68,879 searches from the browsing data of 900 United States adults in March 2025. When an AI summary appeared, users clicked a traditional search result on 8 percent of visits, against 15 percent without a summary. They clicked a source cited inside the summary on 1 percent of visits, and ended their browsing session on 26 percent of pages with a summary against 16 percent without.[12]

Clickstream analysis puts the share of Google searches ending without any click at 69 percent in May 2025, up from 56 percent a year earlier.[13] For business-to-business specifically, Bain found click-through rates on software queries down by as much as 30 percent since AI answers launched, and that 85 percent of B2B buyers purchase from the vendor list they already had in mind before searching.[14] Gartner has separately forecast a 25 percent decline in traditional search engine volume by 2026, though that is a projection rather than a measurement.[15]

> **WHAT THIS DOES TO DEMAND GENERATION**
>
> If 85 percent of buyers buy from a list they arrived with, then being discoverable at the moment of search matters less than being recommended before it. That is a direct argument for advocacy infrastructure rather than more top-of-funnel spend.

3.3 Targeting and measurement have both degraded

One correction is worth making, because the opposite is often asserted. Google did not remove third-party cookies from Chrome. It abandoned deprecation in July 2024 and confirmed in April 2025 that it would not proceed.[16] Cookies remain in Chrome. They are, however, blocked by default in Safari and Firefox, and the practical effectiveness of cookie-based targeting continues to erode.

The better-evidenced disruption is Apple's App Tracking Transparency. Research hosted by the Federal Trade Commission found the share of trackable United States Apple traffic fell 55 percentage points, from 73 percent to 18 percent, that trackable impressions command roughly 51 percent higher prices, and that publisher advertising revenue from Apple users fell around 21 percent. The same work notes that conversion attribution became materially harder, because the replacement reporting operates at campaign rather than user level.[17] A UCLA Anderson working paper measured the cost side directly: removing offsite tracking data raised the median cost per incremental customer by 37 percent, from $112.69 to $154.77, with the largest effects on small businesses.[18]

3.4 Which promotes the base to a growth channel

The consequence of the three trends together is a reallocation. Expansion revenue from existing customers now represents roughly 40 percent of total new annual recurring revenue at the median B2B software company, and exceeds 50 percent for companies above $50 million in revenue.[19] Median net revenue retention, meanwhile, has compressed to around 101 percent from roughly 105 percent in 2021, with wide variance by segment: about 118 percent for enterprise accounts against about 97 percent for small business.[19]

Read together, those two facts say something specific. The existing base is where most new revenue comes from, and the efficiency of getting it has been falling. That is the definition of a growth channel that has outgrown its infrastructure.

ON TWO FAMILIAR STATISTICS

Two figures dominate this conversation and both are routinely misattributed. The claim that a five percent improvement in retention raises profit by 25 to 95 percent originates with Reichheld and Sasser in Harvard Business Review in 1990, where the measured range was 25 to 85 percent across specific industries. The wider 25 to 95 percent range is a later restatement. The related claim that acquiring a customer costs five to 25 times more than retaining one is a rule of thumb popularised in Harvard Business Review in 2014, not a controlled finding. We cite both here as framing, not as evidence.

4 The Customer Growth Maturity Model

There is no dominant analyst framework for customer-led growth treated as a single discipline spanning referral, advocacy, loyalty and lifecycle. The closest adjacent work is Forrester's Customer Advocacy Model, published in May 2025, which stages the build-out of a proactive advocacy program, and Forrester's Customer-Led Sales Growth Model.[20] The model below is Ambassador's, and it differs in scope: it treats data unification and attribution as maturity axes in their own right rather than as reporting concerns.

LEVEL	MODE	WHAT IT LOOKS LIKE IN PRACTICE
L1	Reactive	Spreadsheets and manual outreach. Referrals happen by accident. No structure, no measurement.

LEVEL	MODE	WHAT IT LOOKS LIKE IN PRACTICE
L2	Point Solution	A single tool serving a single audience. Data siloed. Attribution absent or manual.
L3	Coordinated	Several tools connected by integration. Partial automation. Attribution still assembled by hand.
L4	Intelligent	One connected operating system. HiroAI actively guiding. Attribution improving across engines.
L5	World-Class	Full attribution across nine engines. CFO-ready reporting. Agent-orchestrated execution. Intelligence compounding.

Given that only 14 percent of organisations report a single customer view,[3] and that stack utilisation sits near half of purchased capability,[6] the population concentration at L1 and L2 is unsurprising. Competing point solutions in this category generally top out at L3, where tools are connected but attribution remains manual.

HiroAI is present from day one at every level. It does not unlock at L5. The operative framing in a discovery conversation is simple: you start where you are, and the orchestrator starts with you.

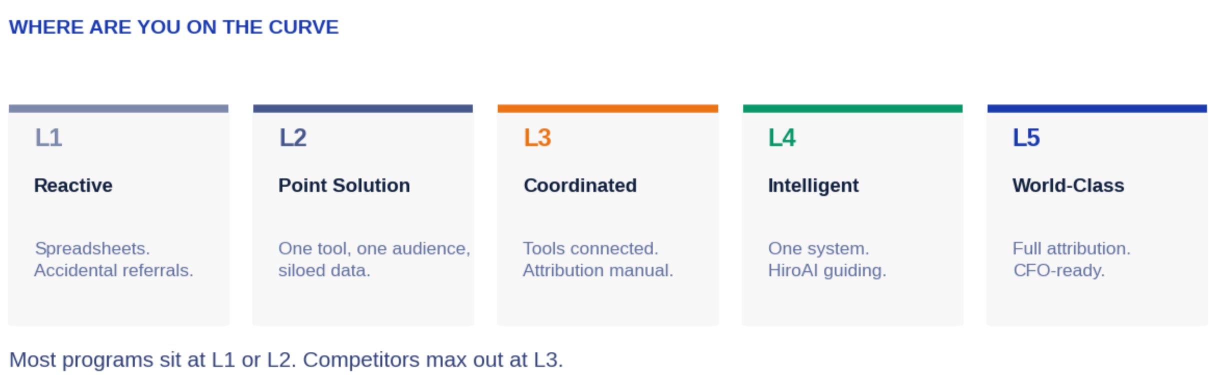

Figure 4. The five levels, at a glance.

4.1 Reading the curve

Plotting both paths against maturity makes the argument concrete. The horizontal axis is lifecycle maturity. The vertical axis is revenue from the customer base already owned. The dotted line is cost of acquisition.

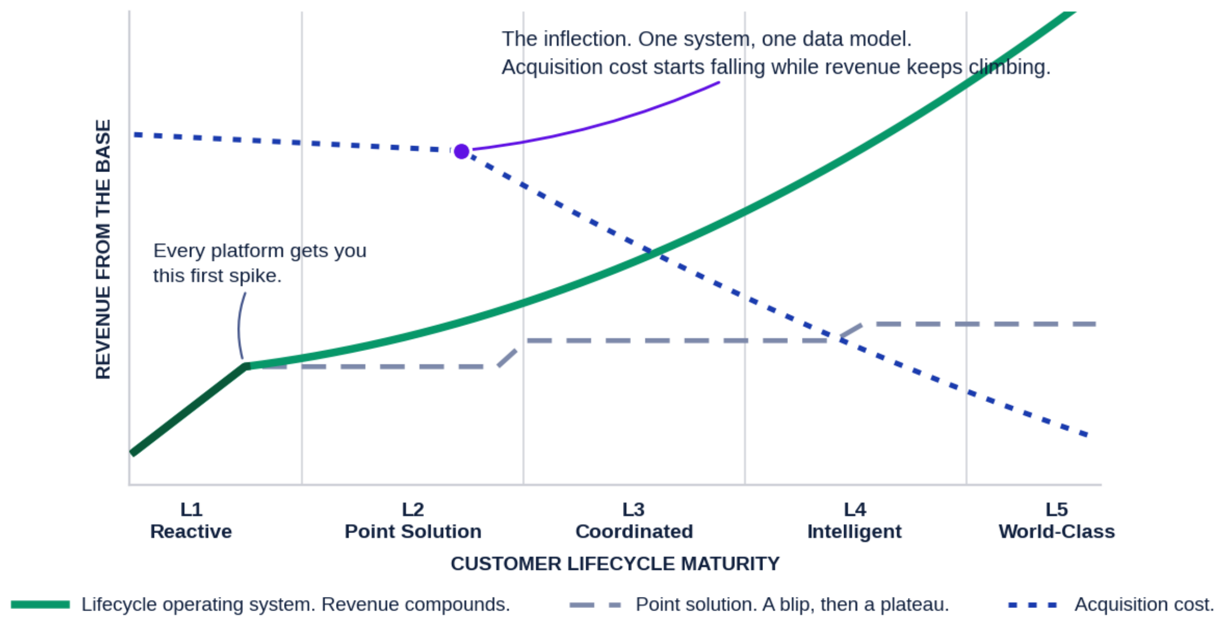

Figure 5. Both paths share the same opening spike. They separate on what happens after it.

Two features of this chart carry the whole argument. First, both paths share the same first spike. Any competent referral tool delivers it, and if that is all a business needs, it should buy the cheapest one that does it well. Second, the paths separate on what happens afterward. The point solution path produces a blip, a plateau, another blip when a feature is added, another plateau. Each blip is real. None compound, because each is a separate system with a separate data model.

Around L3 the dotted line breaks downward. That inflection is the reason this argument belongs in front of a chief financial officer rather than a program owner.

5 Why the second audience requires one system

A company can buy a second point solution for the second audience. Many do. It does not produce the compounding curve, and the reason is mechanical rather than ideological.

Advocacy, product, partners, lifecycle signals, and paid media all generate signals about the same person. Held in five systems, they yield five partial pictures and no way to act on the whole. A referral invitation cannot be triggered by a survey response if the survey tool and the referral tool do not share a customer record. Someone mid-referral cannot be retargeted if the media platform has never heard of the referral funnel. And none of it can be attributed, because attribution requires a single spine, which is precisely the capability the signal-loss research in section 3.3 shows has become scarcer.[17]

WHY THE SECOND AUDIENCE ONLY WORKS ON ONE SYSTEM

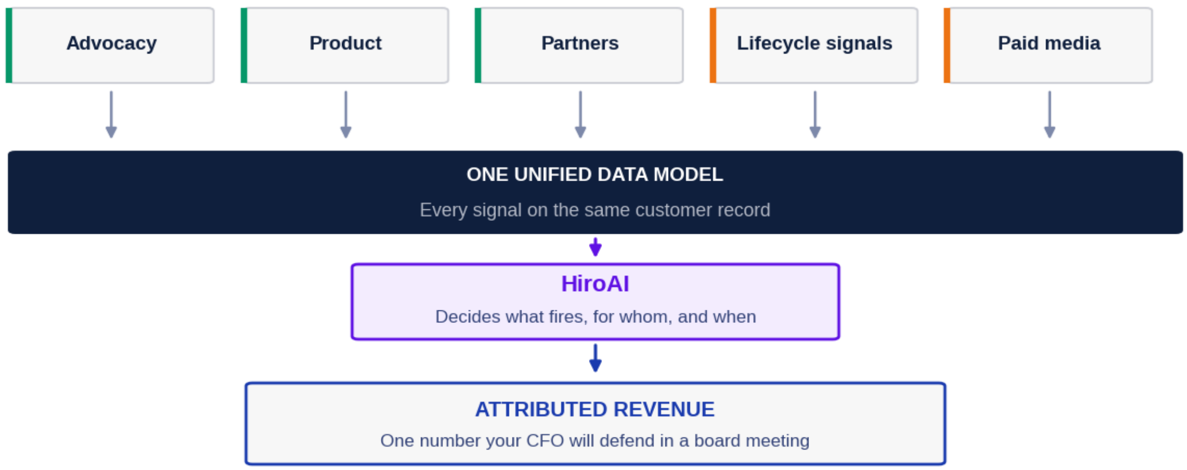

A point solution can hold one of the boxes on the top row. It cannot hold the row.

Figure 6. One unified data model, one orchestrator, one attributable output.

One unified data model places every signal on the same record. An orchestrator above it decides what fires, for whom, and when. Because every action and every dollar traces through one system, the output is a number a chief financial officer will defend rather than tolerate.

> *A point solution can hold one of those boxes very well. It cannot hold the row.*

5.1 What the architecture actually is

Ambassador is The Customer Lifecycle Operating System. One orchestrator, nine engines, and seven specialist agents on a unified data model. The engines are organised into three outcomes.

PILLAR	OUTCOME	ENGINES
GROW	Acquire through advocacy	Advocacy, Prospect, Programmatic
KEEP	Retain through lifecycle	Retention, Incentive, Communication
PROVE	Prove the revenue	Predictive, Attribution, Finance

HiroAI is the orchestrator at the centre, not a chatbot and not an add-on. It runs the nine engines as one closed loop and coordinates the agent layer that executes inside them.

6 The economics of the inflection

Growth leaders are asked two questions in every budget cycle. Is this producing revenue, and is it producing revenue efficiently. Most customer program spend can answer the first and not the second.

The compounding equation. Cost down, value up, at the same time.

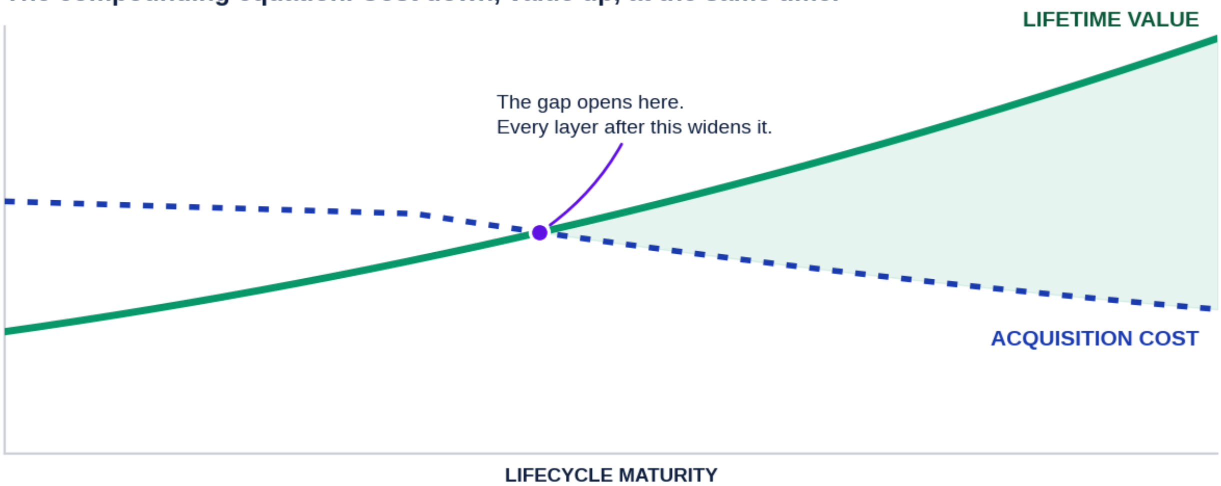

Figure 7. Lifetime value rises while acquisition cost falls, for the same underlying reason.

Across the maturity curve, lifetime value climbs while acquisition cost falls simultaneously. Each additional advocacy audience is cheaper to activate than the last, because the infrastructure exists and the intelligence has accumulated. Each additional lifecycle signal sharpens targeting, which raises conversion, which lowers cost per acquisition again. The retention side of that equation is the durable one, consistent with the finding that the referral retention advantage persists while the margin advantage decays.[1]

The gap between the two lines is the business case. It does not open at the first spike. It opens at the inflection, and every layer added after it widens the gap.

6.1 The argument that wins the budget

One further point tends to decide these conversations at executive level, and it is not about revenue at all.

A point solution puts the entire cost inside one program owner's budget, and grows that line every year the program works. It is a strange penalty for success, and every program owner who has been through a renewal recognises it. An operating system spreads the same spend across marketing, sales, retention, and product, because each of those functions receives something in return. Marketing gets acquisition. Sales gets pipeline. Retention gets save motions and reward infrastructure. Product gets in-product growth surfaces.

THE CFO FRAMING

Same revenue. Diversified operating expense. That is a materially easier number to defend, and it is usually why a conversation that began with a program manager ends up in front of a chief revenue officer.

7 The path, staged

Nothing in this paper argues for buying a whole lifecycle at once. It argues against it. The sequence below is the one that works, and each stage is independently defensible.

THE PATH

Advocacy is the entry point. It is not the ceiling.

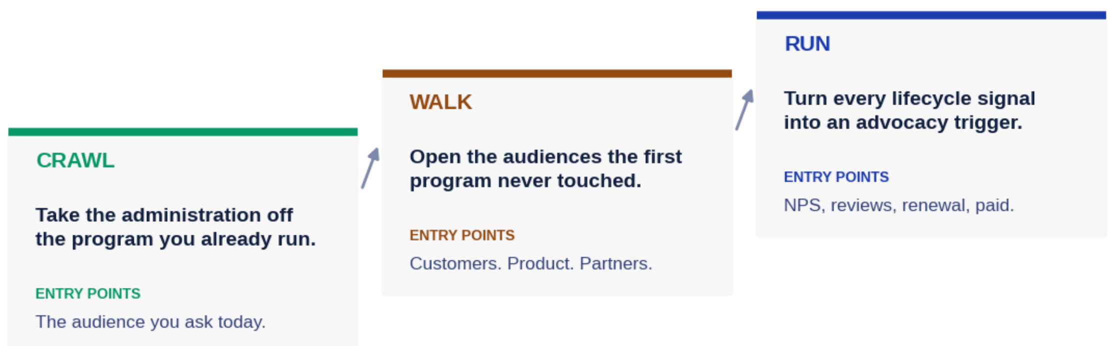

Nobody buys the whole lifecycle on day one. Each layer makes the one before it worth more.

Figure 8. Three stages, each making the last worth more.

Crawl. Take the administration off what already runs

Same audience, same structure, none of the manual work. Rewards fire on the CRM stage chosen rather than on a monthly report. Every participant gets a live leaderboard and a personal dashboard. Entry is a link or QR code with no login, because every login required is participation lost. Payouts reconcile and ledger themselves before finance sees them.

The objective at this stage is not new revenue. It is removing the person holding the program together, so that the next stage becomes possible at all.

Walk. Open an audience the first program never touched

Customers, if the start was internal. The product surface, if the start was people. Partners, if the start was customers. Let them feed each other: the customer referred last quarter becomes the referrer this quarter. This is where the curve resumes climbing, and it is typically the cheapest incremental revenue available to the business. Given that expansion already accounts for around 40 percent of new revenue at the median,[19] it is also the audience most likely to be under-instrumented relative to its contribution.

Run. Turn every lifecycle signal into an advocacy trigger

Surveys, NPS, reviews, renewals, and retention events become moments where the ask is timed to enthusiasm rather than to a campaign calendar. Programmatic follows the people already inside the referral funnel who have not finished. Attribution ties every dollar back to the signal that started it.

8 A diagnostic before any vendor conversation

Three questions, worth answering internally before evaluating anyone, including us.

→ **What did the last program lift, and has it flattened?** If it is still climbing, the business is early and should keep going. If it flattened, the first audience's edge has already been found.

→ **Which audiences could refer you and are not being asked?** Write them down. Customers, employees, partners, in-product users, high-scoring survey respondents. The list is almost always longer than the program.

→ **Can any of it be attributed to revenue?** Not activity. Revenue. If the answer requires someone to build a spreadsheet, the problem is the data model, not the reporting.

CONCLUSION

Flattened, more than two unasked audiences, and no attribution is the exact profile this paper describes. It is also the profile where a staged move produces the largest and fastest return, because the channel has already proven itself and only the architecture is holding it back.

Ambassador is The Customer Lifecycle Operating System. One system that keeps advocacy running and reaches everywhere else it should be firing, orchestrated by HiroAI. If the question is where a business sits on this curve, that is a 45 minute conversation rather than a demo.

REFERENCES

Sources and source quality

Peer-reviewed and primary sources are marked. Commercial and vendor-produced research is labelled as such, because it is directional rather than independently audited. Where a widely repeated statistic misstates its original study, the original is given.

1 Schmitt, P., Skiera, B., and Van den Bulte, C. (2011). **Referral Programs and Customer Value.** Journal of Marketing, 75(1), 46 to 59. DOI 10.1509/jm.75.1.46. Managerial summary published as Do Referral Programs Increase Profits? GfK Marketing Intelligence Review (2013). *Peer-reviewed. Approximately 10,000 accounts, 33-month observation window.*

2 Nielsen (2021). **Trust in Advertising Study.** Fielded August to September 2021, more than 40,000 consumers. *Commercial research. The 88 percent figure is the global result. Earlier Nielsen waves report different figures and should be dated individually.*

3 Gartner (2021). **Cross-Functional Customer Data Survey.** *Named analyst source.*

4 CDP Institute. **The Mass Martech Modernization.** Survey of 200 brands, United States and United Kingdom. *Industry body, survey-based.*

5 Brinker, S. and Riemersma, F. (2024). **State of Martech 2024** and the Marketing Technology Landscape. chiefmartec and MartechTribe. *Recognised industry census.*

6 Gartner (2023). **Marketing Technology Survey**, 405 marketing leaders, May to June 2023; and Gartner (2025) Marketing Technology Survey. *Named analyst source. Utilisation figures differ by survey year and should be cited with the year attached.*

7 Benchmarkit (2025). **SaaS Performance Metrics Benchmarks.** Approximately 1,000 B2B software companies. *Reputable industry benchmark.*

8 Aleph and Benchmarkit (2026). **SaaS and AI Performance Benchmarks.** 342 companies, fiscal 2025 actuals. *Reputable industry benchmark.*

9 Meta Platforms, Inc. **Fourth Quarter and Full Year 2025 Results.** Filed with the United States Securities and Exchange Commission. *Primary source.*

10 WordStream (2025). **Google Ads Benchmarks.** More than 16,000 campaigns, April 2024 to March 2025. *Vendor research.*

11 Paddle and ProfitWell (2025). **Customer acquisition cost trend analysis.** *Vendor estimate. Not an independently audited time series. Corroborating analysis of the same underlying survey data reports approximately 65 percent over five years.*

12 Pew Research Center (22 July 2025). **Google users are less likely to click on links when an AI summary appears in the results.** Browsing data from 900 United States adults, 68,879 searches, March 2025. *Primary research with published methodology. The strongest source in this section.*

13 Similarweb (July 2025), zero-click search analysis; and SparkToro (2026), analysis of Similarweb clickstream data. *Commercial clickstream research. Figures move quarterly and should be date-stamped in use.*

14 Bain and Company (September 2025). **Losing Control: How Zero-Click Search Affects B2B Marketers.** Companion consumer survey conducted with Dynata, December 2024, more than 1,100 United States consumers. *Named analyst source, business-to-business specific.*

15 Gartner (February 2024). **Forecast on search engine volume.** *Analyst projection, not a measurement. Presented in this paper as a forecast.*

16 Google (July 2024 and April 2025). **Announcements on Privacy Sandbox and third-party cookies in Chrome.** *Primary source. Google did not proceed with deprecation. Statements that third-party cookies have been removed from Chrome are inaccurate.*

17 Skiera, B. et al. **Research on the effects of Apple App Tracking Transparency.** Hosted by the United States Federal Trade Commission. *Academic research.*

18 UCLA Anderson School of Management (May 2025). **Working paper on the advertising cost effects of removing offsite tracking data.** *Academic working paper, not yet final publication.*

19 Benchmarkit and Maxio. **Net revenue retention and expansion revenue benchmarks, 2021 to 2024.** *Reputable industry benchmarks. Private B2B software companies.*

20 Forrester (6 May 2025). **Introducing The Forrester Customer Advocacy Model**; and Forrester, **Customer-Led Sales Growth Model.** *Named analyst frameworks. Referenced here to position the Customer Growth Maturity Model, which is Ambassador's own framework.*

21 Reichheld, F. F. and Sasser, W. E. Jr. (1990). **Zero Defections: Quality Comes to Services.** Harvard Business Review, 68(5), 105 to 111. *Primary source for the retention and profit relationship. The measured range was 25 to 85 percent across specific industries. The commonly quoted 25 to 95 percent range is a later restatement associated with The Loyalty Effect (1996).*

22 Gallo, A. (29 October 2014). **The Value of Keeping the Right Customers.** Harvard Business Review. *Popularised the claim that acquiring a customer costs five to 25 times more than retaining one. This ratio is a rule of thumb aggregated across studies, not a controlled finding, and is presented in this paper as framing only.*

NOTES

Ambassador platform figures cited in the executive summary are Ambassador's approved external proof figures. The Customer Growth Maturity Model is Ambassador's own framework. Benchmark figures drawn from commercial research reflect the populations surveyed by their publishers and are not directly comparable across sources, since definitions of acquisition cost and retention differ by publisher. Search and advertising figures are current as of the dates given and move quickly. Forward-looking capabilities named here, including Surveys, NPS, Reviews, Programmatic Audiences, and AI Agents, are planned and subject to availability and then-current terms. Nothing in this document commits Ambassador to a release date.