A WHITE PAPER / FROM THE AUTHOR OF THE FORTHCOMING *We're Doing This Wrong* № 01 · JUNE 2026

Lost *at* Scale.

Why your startup engagement isn't really engaging anyone—and how to rebuild it for long-term value.

DC Cahalane

Founder, investor, and entrepreneur ecosystem builder.
Author of **We're Doing This Wrong: How To Build Entrepreneur Ecosystems That Actually Work**.

CP CANTILLION PRESS

THE ARGUMENT, COMPRESSED

The largest startup engagement programs in the world have, almost without exception, *stopped engaging anyone.*

LOST AT SCALE · FROM THE OPENING

ABOUT THIS PAPER

Most corporate "for Startups" programs were designed to build relationships and have drifted into distributing discounts. *Lost at Scale* is a forty-minute argument for what was lost, what it costs, and what a rebuilt model — measured over years rather than quarters — actually looks like.

Written for CEOs, CFOs, CMOs, CHROs and program leads inside the major technology multinationals. Built around the cohort economics, the founder perspective, and a practical seven-pillar framework for moving from *for Startups* to *Company × Startup*.

WHAT'S INSIDE

DC

DC Cahalane

Ecosystem architect · four-time founder · VC-in-residence at UCC.
Author of the forthcoming *We're Doing This Wrong*.

CP CANTILLION PRESS

WP·LOST·01
cantillion.press

CONTENTS

Lost at Scale

DUBLIN · 11:47 PM · A TUESDAY

Maya is at her kitchen table.

Seventeen tabs open. Three weeks into building her company. She needs to pick a cloud infrastructure provider by Friday. The startup program landing page of one of the major vendors is glowing in front of her.

The form asks for fifteen things she does not have — company registration, twelve-month ARR projection, tax ID, deck, traction paragraph. Her tea has gone cold. Her co-founder is asleep upstairs. Two emails from an angel investor sit unread because she does not yet know how to answer them.

She fills in what she can. She submits.

The auto-reply lands within seconds. *A member of our team will be in touch within ten business days.* She closes the laptop. She has, in some sense, just made a decision that will shape the next ten years of her company. She does not know that yet.

Three weeks later, she has heard nothing.

The largest startup engagement programs in the world have, almost without exception, stopped engaging anyone.

They process founders through portals, distribute discounts in exchange for places in their funnel, and report the results on quarterly pipeline dashboards. Whatever they once were, they are no longer relationships. They are transactions, dressed in the language of partnership.

Most of the people running those programs see the drift clearly. They are losing the fight to reverse it from inside their own org charts. The argument that follows is addressed to the people above them.

> **This paper is about how that happened, what it costs, and what it would take to build something better in its place.**

40 min	7 + 1	2
End-to-end read	Sections, plus a practical seven-pillar framework	Composite founders threading every page

PREFACE

About this paper

This paper is written for the people inside large technology multinationals who run, fund, or commission what have come to be known across the industry as "for Startups" programs — the engagement initiatives operated under that family of names by virtually every major enterprise software, cloud, and platform company.

It is not a critique of the people running those programs. Many of them are doing excellent work under structural conditions that make excellence nearly impossible. It is a critique of the model those programs have drifted into — and an argument for what a better model looks like.

The argument rests on a simple observation: most startup engagement programs today are optimised for the wrong thing. They are measured as discount-distribution mechanisms with a sales-pipeline attached. They were designed, originally, as something else entirely — something that created durable commercial value for both sides. That original purpose has been quietly lost as programs grew and were absorbed into commercial organisations with commercial KPIs.

The cost of this drift is not abstract. It is measurable, in churn, in brand damage, in lost strategic intelligence, and in long-term pipeline left on the table. Founders who felt genuinely supported in their earliest days become extraordinarily loyal customers as they scale. Founders who were filtered through a portal and sold to on a webinar do not.

This paper lays out the problem, the cost, and a practical framework for moving from the "for Startups" era to something different. The naming matters. "for Startups" positions the founder as a recipient of corporate generosity — a junior partner gratefully accepting a discount in exchange for a place in the funnel. The model proposed here borrows its framing from somewhere more contemporary: the **"Company × Startup"** model, in the manner of brand collaborations and platform partnerships now common across consumer culture, where two parties bring something distinct, both contribute meaningfully, and value is understood to flow in both directions. The shift in language is not cosmetic. It is the whole argument, compressed.

The paper draws on twenty-five years of observing and working inside startup ecosystems across Ireland, the UK, and the United States, and on direct conversations with founders who have been on both sides of the corporate–startup relationship.

You will notice that two founder stories run through the paper. They are composites, drawn from dozens of real experiences, and they are here because the difference between a well-run startup engagement and a poorly-run one is most clearly felt, not analysed. Their journeys are meant to help different readers — the CFO funding these programs, the CMO overseeing them, the CHRO considering their retention implications, the program lead running them day-to-day — each find themselves in the narrative.

READER'S GUIDE

How to read this paper

This paper is written for multi-stakeholder corporate audiences. The argument runs end-to-end, but different readers will care about different parts. If you are sharing this internally, the pointers below may help your colleagues find the section that matters to them.

QUICK REFERENCE · WHERE TO START, BY ROLE

CEOs & CTOs	Section 4 (the economic case) and Section 7 (why this matters now).
CFOs	Section 4, including the cohort model and methodology note.
CHROs & L&D leads	Section 5, Hidden Cost Two — the cultural and retention costs — and Section 6, Pillar 3 (the internal mentor program).
CMOs & brand leads	Section 3 (the founder perspective) and Section 6, Pillar 7 (tone and language).
Program leads	Section 6 in full, including the seven pillars and the twelve-month roadmap.
Limited time	The Executive Summary covers the argument in roughly two minutes.

The paper is approximately a forty-minute read end-to-end. The narrative thread following Maya and Daniel — two composite founders representing the old and new models of engagement — runs through every section as a sidebar. A reader can follow just the founder narratives in roughly five minutes and still come away with the core of the argument.

Executive summary

THE PROBLEM

Corporate startup programs have drifted from their original purpose. What began as relationship-driven, ecosystem-contributing work has been absorbed into sales and marketing organisations and restructured around quarterly pipeline. The result is a transactional, portal-driven, discount-focused model.

WHAT HAS BEEN LOST

Three things: the human touch that made founders want to engage; the ecosystem positioning that made these companies part of how a region's startup community actually worked; and the earliest, most valuable window in a founder's journey — when she is deciding which stack to build on. Most programs arrive after that decision. A discount on a choice already made is not engagement. It is a rebate.

WHAT IS AT STAKE

The opportunity cost is substantial. Founders who get genuine, useful engagement become extraordinarily loyal long-term customers; winning a founder in year one and retaining her through Series C dwarfs the cost of meaningful early support. Beyond revenue: real-time intelligence on emerging technology, internal cultural benefits, and brand position no marketing budget can replicate.

WHAT THIS PAPER PROPOSES

A practical model built on seven ideas: a founder-minded program lead; a compensated startup advisory board; opt-in internal mentorship; a map of genuine corporate value rather than a menu of discounts; cross-functional budget ownership; metrics measured over years, not quarters; and a shift in tone from selling to serving. These are not radical ideas — they are what the best early programs already did.

01

SECTION

The origins

How startup engagement programs actually began — and why remembering matters now.

§ 01 / 07   APPROX. 6 MIN READ

FOUNDER NARRATIVE

Two founders, three weeks in.

Maya Chen

Ledgerwise · AI fintech, pre-seed

It is 11:47 on a Tuesday night and Maya is still at her kitchen table in Dublin. She has been building Ledgerwise — an AI reconciliation platform for small and medium-sized enterprises — for just under three weeks. Her co-founder is asleep. Her laptop has seventeen tabs open. She needs to pick a cloud infrastructure provider by Friday.

The startup program landing page of one of the major cloud vendors is glowing in front of her. The application form asks for her company registration number, her projected twelve-month ARR, her funding stage, her tax identification, her team size, her pitch deck, and a paragraph describing her traction. She has almost none of these. She fills in what she can. She submits. She receives an automated email: **"Thank you for your application. A member of our team will be in touch within ten business days."**

Three weeks later, she has heard nothing.

Daniel Okafor

Stacklane · AI productivity, pre-seed

Daniel is also three weeks in. Stacklane is an AI productivity platform for distributed engineering teams, and he is building it from a co-working space in Manchester. Last Thursday he went to a small founder meetup organised by a local ecosystem group. Standing near the coffee table, he got into a conversation with a woman named Priya, who turned out to run the startup engagement team at one of the major productivity software companies.

She asked him what he was building. She asked him where he was stuck. She gave him her business card and said, **"I'd genuinely love to hear more. Send me an email next week."** He did.

Yesterday morning, Priya replied. No form. No portal. Just a short, warm email: **"Great to meet you last week. Here are a few ways I think we might be able to help."**

THREE WEEKS OF COMPANY-BUILDING

where each founder's vendor relationship actually sits

	Maya LEDGERWISE	VS **Daniel** STACKLANE
FIRST TOUCHPOINT	Landing page · **21 days ago**	Founder meetup · **7 days ago**
WHO INITIATED	Maya did, via the application form	Priya did, in person
REPLY RECEIVED	One auto-reply · "within 10 business days"	One personal email · **under 24 hours**
ANYONE AT THE VENDOR KNOWS WHAT THEY'RE BUILDING	No one	Priya does · she asked
HOW CREDITS SURFACED	Foregrounded on the marketing page	Not raised
NAMED HUMAN ON THE FOUNDER'S SIDE	None	**Priya Anand**

Before the portal

There is a tendency, when looking at any mature corporate function, to assume that its current form is its natural form. The startup engagement programs run today by the major technology multinationals — largely online, application-gated, discount-centric, sales-pipeline-accountable — feel like the obvious way these things would work. They are not. They are, in most cases, the third or fourth iteration of a model that started out looking very different.

Ten and fifteen years ago, most large technology companies engaged with startups through people. A named individual, often someone with founder experience, would show up at ecosystem events, know the local accelerators and investors by name, and maintain direct, long-running relationships with founders in their patch. They might hand out free credits or discounts, but those credits were the end of a conversation, not the beginning. The conversation itself — the mentorship, the introductions, the technical advice, the occasional nudge into a procurement conversation at the right moment — was the product.

These programs were often small. They were frequently owned by business development, partnerships, or even the CTO's office rather than by sales or marketing. Their budgets were modest, their headcount was lean, and their success was assessed qualitatively as much as quantitatively. The program lead was, in effect, an ecosystem ambassador — judged on the strength of the relationships they had built and the positioning their company enjoyed in the startup community, as much as on any quarterly number.

It is worth pausing on how unusual this looks by the standards of contemporary corporate accountability. A program lead with a $2M annual budget and no quarterly pipeline target would, in most companies today, be considered an indulgence — an unjustifiable line item that finance would zero out at the first opportunity. A decade ago, those same finance teams accepted the ambiguity, because the people inside the room could see what the program produced even when the dashboard could not. The CIO who took a meeting because a startup program lead had vouched for the founder. The acquisition that closed three years after a relationship started at an unrelated demo day. The procurement decision that went a particular way because, when the buyer rang around her network, the answers about that vendor came back warmer than they would have otherwise.

The lineage of these early programs is worth tracing too. Most of them did not start inside the corporate parent. They started as something a single committed individual — often a former founder herself, often hired from inside an accelerator or seed fund — argued into existence by force of personality, with sympathetic executive cover and a budget assembled out of three or four functional pockets. The first generation of program leads were, in many cases, building their own job description as they went. The institutional formality came later. The texture of the work, and the texture of the relationships, came first.

> The credits were a gift within an existing relationship, not an acquisition mechanism.

What those early programs actually did

It is worth being precise about what the early, working version of corporate startup engagement looked like, because the model is often remembered more fondly than it deserves and the specifics are worth recovering. The nostalgia for these programs sometimes outruns the evidence; it is important to separate what they actually did from what observers, in retrospect, wish they had done.

The best early programs did five things consistently, and the combination — not any single tactic — is what made the model coherent.

01

Showed up in ecosystems

Their people attended demo days, judged pitch competitions, spoke at founder events, and were physically present in the places where founders were.

02

Invested in intermediaries

Accelerators, incubators, VCs, and angel groups — the trusted filters through which founders actually discovered corporate partners.

03

Offered real help

Not just credits, but technical architecture reviews, customer introductions, GTM advice from people who had actually shipped, and occasional co-marketing.

04

Played a long game

Program leads with multi-year tenures, founders who knew them by name, and institutional memory about who had been engaged at what stage and why.

05

Contributed to the ecosystem

They added to the commons they operated in — sponsoring meetups, paying speakers, funding programmes — rather than extracting from it.

Each of these behaviours is, on its face, modest. None of them is the kind of thing a board strategy day produces. But the five reinforce each other in specific ways. **Showing up without offering real help is tourism. Offering help without intermediary relationships** becomes either patronage or favouritism. **Investing in intermediaries without playing the long game** produces transactional referral flows that decay within a year. **Playing the long game without contributing back** corrodes the goodwill that made the relationships possible in the first place. The reinforcement, not any single tactic, is what produced the outcomes the model became known for.

Crucially, the discount or free-credit element was almost always framed as a consequence of engagement, not the reason for it. A founder who had built a genuine relationship with the corporate partner would, at some point, receive a nudge: **"by the way, you should make sure you're taking advantage of the credits, they'll help with your runway."** The credits were a gift within an existing relationship, not an acquisition mechanism. The sequencing matters more than it might first appear: when the credit conversation comes second, the founder reads it as generosity. When it comes first, she reads it as bait — even when the underlying offer is identical.

Why the model worked

The original model worked for a specific reason: it aligned the incentives of the program lead with the incentives of the founder. A program lead whose success was measured on ecosystem positioning and long-term relationships would naturally invest time in founders who were not yet customers, were not yet ready to be customers, and might not be customers for years. They would help a founder solve a problem that had nothing to do with their employer's product, because doing so built trust, and trust compounded.

This alignment is not an abstract design principle. It is the single most important thing to understand about why the early model produced the outcomes it did. Every meaningful innovation in how large companies engage with smaller ones — from the early days of Bell Labs working with university spinouts, to the IBM partner programs of the 1980s, to the cloud-platform startup initiatives of the 2010s — has, at its best, rested on the same insight: that the most durable commercial relationships are built when the senior party in the relationship has been given permission, by their own organisation, to not press for short-term advantage. Remove that permission and the relationship reverts to a sales interaction. Founders read the difference within a single conversation.

Founders, in turn, treated those program leads as genuine allies. They returned their calls. They sent them term sheets to review. They invited them to board meetings as informal advisors. They mentioned their employer's products in conversations with other founders, not because they had been asked to, but because they actually meant it. The commercial outcomes — infrastructure spend, enterprise licences, eventual partnerships, occasional acquisitions — followed as a second-order effect of the relationship.

It is worth being concrete about how that second-order effect actually showed up in revenue. A program lead who had built a real relationship with a founder at pre-seed would, two or three years later, receive a quiet note when that founder was making a stack decision. The note rarely said **"please send me a quote."** It typically said something closer to **"I'm comparing options, and I wanted to make sure you knew."** The founder was not asking for a discount. She was giving the incumbent a chance to compete on merits, because the incumbent had earned that chance. Those moments — quiet, undramatic, often invisible to the wider sales organisation — were where the long-term value of the program actually accrued. The pipeline did not show them. The relationship lead knew about every one of them.

The model also worked because it was staffed by the right people. The best program leads in this era tended to have founder backgrounds themselves, or to have worked closely enough with founders — at accelerators, in early-stage VC, in angel investing — that they understood the texture of early company-building. They knew what it felt like to be at Maya's kitchen table at 11:47 on a Tuesday night, seventeen tabs open, too tired to fill in a seven-field form. That lived understanding shaped every interaction they had. It also shaped what they would and would not ask a founder for. A program lead with founder context, asked by her own sales organisation to push a startup into an awkward upsell conversation, knew how to refuse — and knew how to make the refusal look like a strategic choice rather than insubordination. Without that context, the same request gets executed, the founder feels it, and the relationship that took two years to build is damaged in a single exchange.

> The commercial outcomes followed as a second-order effect of the relationship.

The seeds of what went wrong

Even in their best early versions, these programs carried within them the seeds of what would go wrong later. The model that worked, worked despite four structural fragilities that were present from the beginning. Each of them, in isolation, was manageable. Together, and once the programs began to scale, they would prove fatal.

FIRST. They were expensive per founder reached. A program lead who could only cover two or three hundred meaningful founder relationships in a year was, by definition, a bottleneck. As the startup ecosystem grew globally — and the 2010s saw an explosion in the sheer number of venture-backed startups being founded each year — pressure mounted to scale engagement beyond what relationship-based models could support. The honest version of this story is that the program leads themselves often saw the limit clearly and lobbied internally for more headcount; the resourcing rarely arrived in proportion, and the gap was eventually closed with technology — portals, application forms, automated nurture — rather than with people.

SECOND. They were hard to measure. A senior executive asked to justify the program's budget could not easily produce a quarterly dashboard. The commercial value was real, but it showed up three or four or five years later, diffused across multiple accounts, and was difficult to attribute cleanly. This made the program vulnerable whenever a new CFO arrived with a sharper ROI lens, or whenever a quarterly miss elsewhere in the business forced a search for budget that could be reallocated without immediate visible damage. Programs whose value compounds slowly are always at risk in organisations whose attention compounds quickly.

THIRD. They sat awkwardly inside corporate org charts. They were not quite sales, not quite marketing, not quite partnerships, not quite product. Every reorganisation risked absorbing them into a function with incompatible KPIs, and most of them were, eventually, absorbed. The org-chart problem was, in retrospect, the most consequential of the three. Budget and measurement could have been argued for. But once a program reported into a function whose incentives were misaligned with its purpose, the daily decisions that shaped the program's character — what to optimise, what to cut, what to celebrate — began to be made by people whose own success depended on a different outcome.

FOURTH. They were under-protected by senior leadership. The best programs typically had one executive sponsor — a CTO, a head of business development, occasionally a CEO — who personally understood the long-horizon logic and was willing to defend it. When that sponsor moved on, retired, or was promoted into a different remit, the program lost the air cover that allowed it to behave differently from the rest of the commercial organisation. The successor rarely refused to support the program outright. More commonly, they simply allowed it to be slowly assimilated into existing structures, and the assimilation, once underway, was difficult to reverse.

The rest of this paper is, in one sense, about what happened next — the absorption, the industrialisation, the drift into portal-and-discount economics. The four fragilities above are not arguments against the original model. They are the design constraints any rebuilt version has to satisfy.

02

SECTION

Lost at scale

How each defensible decision, compounded over years, produced a model that looks nothing like the one it started as.

§ 02 / 07 APPROX. 7 MIN READ

Four weeks in.

Maya Chen

Ledgerwise · pre-seed

Maya has finally been approved. Eighteen days after submitting her application, she receives a templated email confirming her acceptance into the startup program. She is entitled to twelve thousand dollars in credits, valid for twelve months, usable against a defined set of products, subject to terms and conditions linked in a footer.

There is no introduction to a human being. There is a link to a resource portal, a monthly webinar calendar, and a support email address. The onboarding call is a recorded video.

Maya is, technically, now a member of this company's startup program. She does not yet know anyone at the company. She suspects she never will.

Daniel Okafor

Stacklane · pre-seed

Daniel has just had his second call with Priya. The first was a half-hour conversation in which she asked about his architecture, his team, his target customer segment, and what was keeping him awake.

The second, a week later, was a warm introduction to a senior sales engineer at her company who had spent two years selling into distributed engineering teams and had offered to spend an hour a month talking Daniel through what good early GTM motion looked like.

At the end of the second call, almost as an afterthought, Priya mentioned, **"Oh, and make sure you're taking advantage of the credits — they'll buy you some runway."**

AFTER FOUR WEEKS

the relationship, audited

	Maya LEDGERWISE	VS	**Daniel** STACKLANE
DAYS SINCE FIRST CONTACT	**18 days** from application to approval		**7 days** from meeting to first call
ONBOARDING FORMAT	Templated email + recorded video		30-min call + warm internal intro
FIRST SUBSTANTIVE CONVERSATION ABOUT	Eligibility & terms and conditions		Architecture, team, target customer
KNOWN HUMANS INSIDE THE VENDOR	**Zero**		**Two** · Priya + a sales engineer
CREDITS FOREGROUNDED BEFORE RELATIONSHIP	Yes · central to all comms		No · mentioned as an aside
WHAT THE FOUNDER FEELS	*A member of a list*		*A person being helped*

How the model drifted

Most corporate startup engagement programs did not fail through a single bad decision. They drifted. Each drift was defensible at the time. Each drift, compounded over several years, produced a model that looked nothing like the one the program was originally designed to run.

The drift is worth tracing carefully, because many of the people responsible for these programs today cannot remember a version that worked differently, and the industry as a whole has collectively forgotten that it once did.

i.

The move under sales or marketing

The first and most consequential drift was organisational. Startup engagement programs, originally sat in ambiguous places on the org chart — partnerships, business development, CTO's office — were gradually pulled into the commercial organisation. Usually into marketing, often into sales, occasionally into a hybrid demand-generation function.

The logic was rarely malicious. As programs grew, they needed budget, and budget flowed more easily from organisations with clear commercial accountability. A startup program producing leads, even imperfectly, fit more neatly into a marketing or sales structure than into a vague partnerships one. The reporting lines were cleaner. The KPIs were legible.

The cost showed up slowly. A program lead once judged on ecosystem positioning was now judged on marketing-qualified leads. The time horizon collapsed from years to quarters. The definition of success narrowed from *strong relationships with the ecosystem* to *pipeline attribution*.

ii.

Industrialisation of the founder intake

Once the program was measured on pipeline, the next drift followed naturally. Relationship-based intake — the program lead meeting founders in person, assessing them case by case — could not produce enough volume to satisfy the new KPIs. So intake was industrialised.

The online application portal arrived. The eligibility criteria were formalised. The approval process moved to an internal operations team or, later, to an automated pipeline with tiered approvals. The founder journey was mapped into a funnel: landing page, application, approval, onboarding email, portal login, credit activation, engagement nurture track, upgrade conversation.

None of this is, in isolation, unreasonable. The problem is specific to this context: the thing being industrialised was a relationship. And relationships, unlike campaigns, do not industrialise well. A founder who is processed through a funnel does not feel engaged. She feels managed.

> A founder who is processed through a funnel does not feel engaged. She feels managed.

---

iii.

The rise of the webinar as substitute

As program intake scaled and the program lead could no longer personally engage with each founder, corporate program teams needed a way to deliver "value" at scale. The answer, across the industry, converged on the webinar. Monthly or weekly sessions, hosted by the company, delivered at the program cohort, covering topics notionally of founder interest.

Founders attended. Most of them attended once. They noticed quickly that the webinars were not, in the main, about them. A session titled *"How startups go to market"* would turn out to be a product demo with case-study veneer. A session titled *"Building scalable infrastructure"* would turn out to be a walkthrough of the vendor's own reference architecture, with the vendor's own services in the centre of every diagram.

The webinar became, in effect, a marketing channel disguised as a content benefit. The inverse relationship between quantity and quality became one of the defining signatures of the drifted model.

iv.

The discount as the product

The final drift was the most subtle and the most damaging. Over time, as the human element of the program faded and the scalable elements grew, the startup's perception of the program shifted. What had once been a relationship with a company that included, among other things, a discount, became a discount with some peripheral marketing activity attached.

The discount was simple, legible, and valuable. Everything else the program offered — webinars, nurture emails, a resource portal, a community Slack — was harder to value, and in most cases valued at approximately zero.

The corporate program team, seeing this, doubled down on the discount as the hook. Marketing creative shifted to lead with credit amounts. Landing pages foregrounded the dollar value. The headline became, in effect, *"Startups: get $X of free credits."* The rest receded into the small print. At this point, the program had become its own final form: a rebate mechanism with a light-touch marketing programme attached.

THE CUMULATIVE SHIFT · TEN YEARS OF SMALL REASONABLE DECISIONS

THEN		NOW
— Ecosystem ambassador with multi-year tenure	→	— Demand-gen manager on an annual cycle
— Intake by hand, founder by founder		— Intake by portal, processed in batches
— Help measured in introductions and conversations		— Help measured in webinar attendance
— Discount offered as a consequence of the relationship		— Discount offered as the front door

Why the drifted model is structurally broken

It is worth being precise about why the drifted model does not work — not in the sense of failing to hit commercial targets, because in many cases it does hit them, but in the sense of failing to capture the strategic value that the original model captured. Three structural problems stand out.

A selection problem.

A program that leads with discounts attracts founders who are looking for discounts. Most of those founders have already decided which technology to build on; they are just looking to reduce the cost of that decision. The program therefore engages almost entirely with founders who were already going to be customers, and fails to engage with the founders who are still deciding. The most valuable window — the one in which corporate engagement could actually influence the technology decision — is systematically missed.

A relationship problem.

A founder whose only contact with the company is a portal, a welcome email, and a series of webinars has no relationship with the company. When she needs something — a technical introduction, a sales-motion conversation, a warm route into a potential customer — she has no one to call. She does not experience the company as a partner. She experiences it as a vendor, and usually a distant one.

An ecosystem problem.

A program run at scale through portals and webinars does not contribute to the ecosystems it operates in. Its program leads do not show up at local demo days. Its relationships with accelerators, VCs, and angel groups decay into transactional referral agreements. The company's presence in the startup community gradually thins from *active participant* to *logo on a sponsorship banner.*

A NOTE ON RESPONSIBILITY

Nothing in this chapter is a criticism of the individuals running these programs today. The people working inside large-company startup teams are some of the most thoughtful, committed, and frustrated operators in the technology industry. Many see the drift clearly. Many are trying, within the constraints of their organisational structure, to push back against it. They are, in most cases, losing that fight, because the drift is structural and cannot be reversed from below.

03

SECTION

The founder perspective

What early-stage founders actually need — and what most programs are structurally unable to give them.

§ 03 / 07   APPROX. 6 MIN READ

What founders actually need

To understand what good corporate–startup engagement looks like, it is necessary to start from the founder's point of view, and specifically from the point of view of a pre-seed or seed-stage founder, because that is where the most valuable engagement window exists and where most corporate programs are weakest.

Pre-seed founders are not, in general, suffering from a shortage of free credits. They are suffering from a shortage of time, a shortage of context, and a shortage of trusted advisors. Every decision they make in the first six months of company-building has compounding consequences that they do not yet have enough experience to fully evaluate. The technology stack they choose will shape their hiring, their architecture, their unit economics, and their eventual exit options. They are making decisions that matter, with information they do not have, on timelines they cannot control.

What a founder in this position actually needs from a corporate partner is short to describe and genuinely difficult to deliver: access to people who have done this before, who are prepared to answer honest questions honestly, who are not trying to sell them anything in the next thirty days, and who will respond when called upon. That is it. That is the product.

> Pre-seed founders are not suffering from a shortage of free credits. They are suffering from a shortage of time, a shortage of context, and a shortage of trusted advisors.

SCARCITY I

Time

Running a company with a team of one to three, building product, selling, hiring, and fundraising simultaneously. Every hour spent on a corporate program is an hour not spent on whether the company survives the next ninety days.

SCARCITY II

Context

Every decision in the first six months has compounding consequences the founder does not yet have enough experience to fully evaluate. Stack, hires, pricing, governance — each shapes the next decade, made with incomplete information.

SCARCITY III

Trusted advisors

People who have done this before, will answer honest questions honestly, are not trying to sell anything in the next thirty days, and will respond when called upon. That is the product a serious program off

The five things founders value most

Across twenty-five years of watching these relationships form and dissolve, five kinds of value consistently matter most to early-stage founders. None of them is a discount.

01 Architectural advice from practitioners

At pre-seed, a founder is often making technology decisions she will live with for years. A well-placed conversation with a senior engineer inside a large vendor — someone who can look at the architecture, ask the right questions, and flag the problems the founder has not yet thought to worry about — is worth more than any amount of documentation. It is also, from the vendor's point of view, extraordinarily cheap to provide.

02 Go-to-market coaching from people who have actually done it

Most founders at pre-seed are technical, and the specific craft of enterprise or SMB sales — what a discovery call actually looks like, how to qualify, how to price, how to structure a pilot — is not something they have done. Large companies employ thousands of people who have done exactly this, and many are willing to spend an hour explaining how it works to a grateful founder.

03 Warm introductions to the right people

A founder's access problem is almost always relational. She needs a specific potential customer, a specific potential hire, a specific potential investor. The corporate partner, operating inside a network of tens of thousands of employees, can frequently solve this in a single email — when the program staff have visibility into that network and a mandate to mobilise it.

04 Genuine product feedback, in both directions

A pre-seed founder is often building something adjacent to a large vendor's platform. A genuine conversation between the founder and a product manager — about where the platform is going, where the gaps are, and where the founder's product might complement rather than compete — is valuable to both sides. The founder gets strategic clarity. The vendor gets early signal.

05 The assurance of being seen

This last one is harder to describe and easier to feel. A founder at pre-seed is operating in considerable isolation. To be taken seriously by a significant company — to be treated as a future partner rather than a future lead — is psychologically substantial. Founders remember who treated them this way. They remember it for years.

What founders actively dislike

It is also worth stating, plainly, what founders find actively counterproductive in most current startup programs. This is not a complaint about polish. It is a list of signals founders read as evidence that a program does not take them seriously.

- The long application form that demands data the founder does not have.
- The approval window measured in weeks.
- The welcome email that lists terms and conditions before it lists people.
- The webinar that turns out to be a product demo.
- Credits with hard time limits and narrow-use conditions, calibrated to maximise vendor conversion rather than founder runway.
- The community Slack that is, in practice, a monologue from the vendor's marketing team.
- The onboarding call that is a recording.
- The single point of contact that is an email alias.
- Follow-up sequences cadenced for the vendor's funnel and not the founder's life.

Any one of these signals, taken on its own, is survivable. In combination, which is how they typically arrive, they produce a strong and lasting impression: this company does not see me as a partner. This company is running a process, and I am an input to it. A founder who forms this impression in her first month will carry it for years. She will tell other founders.

The time-poverty problem

The pre-seed founder is, operationally, one of the most time-poor people on earth. Every hour she spends on a corporate program — filling in a form, sitting on a webinar, navigating a terms-and-conditions document — is an hour not spent on the work that will determine whether her company survives the next ninety days.

This means the bar for corporate programs is higher than most program designers realise. A program that demands thirty minutes of a founder's time must repay that thirty minutes with something genuinely worth more than thirty minutes. A program that wants to attract serious founders has to be worth serious founders' time. That means front-loading the value, minimising the admin, and making every single touchpoint repay the founder for the attention she is giving.

04

SECTION

The economics of early loyalty

Why startups are the best customers large companies don't realise they have — and what the numbers actually look like.

§ 04 / 07   APPROX. 8 MIN READ

Eighteen months on.

Maya Chen

`Ledgerwise` · seed, $3M

Ledgerwise has just closed a $3M seed round. Revenue is growing. The team is now eleven people. Maya's infrastructure spend has tripled, and she is beginning to worry about cost at scale.

She has not heard from the startup program at her cloud vendor in seven months. When a competing vendor's account executive cold-emails her asking about a migration, she replies within the hour. A demo is scheduled for the following week.

She does not feel any loyalty that would make the conversation uncomfortable. She is simply evaluating options, as any reasonable CEO would.

Daniel Okafor

`Stacklane` · seed, $4M

Stacklane has just closed a $4M seed round. Daniel has increased his spend with Priya's company by a factor of four over the past year.

When a competing vendor's account executive cold-emails him, he forwards the email to Priya with a short note: **"Anything I should know before I take this call?"** She replies within the hour with context on the competitor's positioning, a candid acknowledgement of where her company is weaker, and an offer to set up a direct conversation with the product team.

Daniel takes the competing vendor's call, but it is a short one. He is not looking to migrate. He is looking, if anything, for reasons to consolidate more of his stack onto the platform he already trusts.

EIGHTEEN MONTHS ON · THE EARLY LOYALTY SIGNAL IS NOW LEGIBLE

	● Maya	● Daniel
Funding stage	Seed · $3M raised	Seed · $4M raised
Spend change vs early	Infrastructure spend ↑ 3×	Spend with vendor ↑ 4×
Last contact from vendor program	Seven months ago	Ongoing dialogue
Reaction to competitor cold email	Replied within the hour; demo next week	Forwarded to Priya with one line
Migration risk	Material — actively evaluating	Low — looking to consolidate
Net dollar retention signal	At risk	Expanding

Startups are the best customers large companies don't realise they have

There is an assumption buried inside most corporate startup programs that treats startups as low-value, high-effort accounts — worth engaging with because some of them will scale, but not themselves particularly valuable in the short term. This assumption is wrong, and it is worth taking apart.

Early-stage startups are, in several specific and measurable ways, among the most economically attractive customer cohorts a large enterprise software company can acquire. They start small. They grow fast. They buy with conviction. They expand their usage aggressively when the product works. They are, in the main, deeply loyal to the vendors who helped them when they were small. And the cost of winning them, at the pre-seed stage, is a tiny fraction of what it will cost to win the same account once the company has reached Series B and the competitive landscape has thickened.

To put this plainly: a founder you help at pre-seed is a founder who may grow into a six-figure annual account by Series A, a seven-figure account by Series C, and an eight-figure account by IPO. The cost of helping her at pre-seed might be a dozen hours of senior employee time and a few thousand dollars in credits. The lifetime revenue difference between winning her early and winning her late — or losing her entirely — routinely runs into the millions. Very few other acquisition channels offer this kind of asymmetry.

> A founder you help at pre-seed is a founder who may grow into a six-figure account by Series A, a seven-figure account by Series C, and an eight-figure account by IPO.

THE COMPOUNDING ACCOUNT · ONE FOUNDER, ONE DECADE

PRE-SEED	SERIES A	SERIES C	IPO & BEYOND
$2–5K	**~$100K**	**$1–5M**	**$10M+**
Credits + senior time. The window where engagement actually shapes the stack.	First paying year. The founder's incumbent bias is set or broken here.	Multi-product consolidation. Decisions taken in year one cash in.	Eight-figure annual relationship. The original engagement repays itself thousands of times over.

The compounding of early loyalty

Founder loyalty is not a soft phenomenon. It is a specific, observable pattern in how startups buy software and infrastructure as they scale. It has four features worth naming.

FIRST. Founders who felt supported early default to incumbent vendors in subsequent purchasing decisions — a CRM, a data warehouse, an identity provider, a new cloud region. The choice is not always conscious; it is the path of least resistance, and trust is what makes the path feel shorter.

SECOND. Loyal founders resist competitive displacement far more effectively than comparable customers acquired through later-stage enterprise sales motions. The cost-to-displace is high precisely because the cost-to-acquire was low.

THIRD. Loyal founders recommend; the startup community is dense, vocal, and networked, and their recommendations carry disproportionate weight that marketing budgets cannot replicate. One generous early intervention will often surface, unprompted, in three or four founder conversations a year.

FOURTH. Loyal founders forgive. Every vendor will, at some point, ship a bad release, fumble a migration, or miss a renewal SLA. Founders who feel known by the vendor absorb those incidents as the cost of doing business. Founders acquired transactionally treat the same incidents as breach-of-trust events — and they tell the network when they leave.

The four features compound. A loyal founder who defaults to the incumbent on her next decision, resists a competitor's displacement attempt, recommends the vendor to peers, and absorbs the occasional fumble without leaving — that single account is producing four distinct streams of value out of one early relationship. None of the four streams shows up on the program's quarterly dashboard. All four are visible in the five-year cohort numbers, and all four are absent from the cohort numbers of programs that did not invest in the early relationship.

This is what is meant, in the cohort table on the next spread, by an 85% retention rate versus a 50% one. The difference is not magic. It is the predictable, structural consequence of having engaged seriously with the founder during the window when serious engagement was possible.

FIGURE 1

Illustrative five-year cohort economics

Two pre-seed cohorts, identical in size and baseline product. The only difference is the depth and nature of engagement.

METRIC	A TRANSACTIONAL COHORT	B RELATIONSHIP COHORT	Δ
Founders engaged (year 1)	1,000	1,000	—
Survive to Series A (5-yr)	150 15%	180 18%	+30
Retained on platform to Series A	75 50%	153 85%	+78
Avg year-5 annual spend	$85K	$135K	+$50K
Cohort year-5 ARR	$6.4M	$20.7M	+$14.3M
Program cost (5-yr)	$3.5M	$6.0M	+$2.5M
Net 5-year return	**$2.9M**	**$14.7M**	**+$11.8M**

ASSUMPTIONS Typical SaaS / cloud-infrastructure pricing; base churn rates from published industry benchmarks; relationship cohort assumes higher retention (85% vs 50%), modestly higher survival, and higher expansion through consolidation onto a trusted vendor.

WHERE THE +$11.8M COMES FROM · DECOMPOSING THE NET RETURN

~70%

RETENTION

Founders who default to the incumbent vendor at Series A rather than re-opening the platform decision.

~22%

EXPANSION

Per-account spend uplift from consolidation onto a trusted vendor as the company grows.

~8%

SURVIVAL

Modest improvement in five-year survival from architectural and GTM mentorship.

Three revenue streams, not one

The customer-spend model captures only one of three distinct value streams a well-run startup program produces. The other two are routinely ignored in ROI conversations, and each of them is material.

I.

Customer revenue

The compounding cohort economics modelled in Figure 1 — survival, retention, expansion. The visible line, and the one that finance teams already track, however imperfectly.

II.

Strategic intelligence

Embedded program leads, an active advisory board, and quarterly intelligence reports deliver early ground-truth on where technology and markets are moving — months before analyst firms catch up. For a strategically-minded CEO or CTO, this is a competitive asset, not a nice-to-have.

III.

Internal culture & retention

Employees who spend a few hours a month mentoring founders consistently report higher job satisfaction and an appetite to stay. The marginal cost of offering this access is negligible; the marginal benefit to engagement and retention is meaningful.

The asymmetric bet

What the economics describe is a radically asymmetric bet. The downside of running a serious, relationship-led startup program is modest: an incremental spend on the order of a few million dollars per year for a global program. The upside, measured across customer revenue, strategic intelligence, and internal culture, routinely runs to tens of millions of dollars per year and produces strategic positioning that competitors cannot replicate.

This is, structurally, the kind of bet that CFOs should love and CEOs should be asking about. The reason it is so rarely made is not that the economics do not work. It is that the current program structure makes the economics invisible. The drifted model hides its own downside and suppresses its own upside; the rebuilt model, described in section six, makes both visible.

METHODOLOGY NOTE

Calibrating the cohort model to your own organisation

Survival rates `15% / 18%`	Draws on published five-year survival data for venture-backed startups, with the relationship-cohort uplift constrained to what GTM and architectural mentorship could plausibly influence at the margin.
Retention rates `50% / 85%`	Reflects the observed difference between commodity-acquired SaaS accounts and high-touch, relationship-led accounts in published industry benchmarks on net revenue retention and gross retention.
Average year-5 spend `$85K / $135K`	Assumes typical SaaS or cloud-infrastructure pricing for a Series A-stage venture-backed company, with the relationship-cohort uplift modelling the net-revenue-retention effect of consolidation onto a trusted vendor.
Program cost `$3.5M / $6.0M`	Assumes a global program covering approximately 1,000 founders. The incremental $2.5M reflects senior staffing, advisory-board compensation, and internal-mentor coordination as described in Section 6.

Substitute company-specific inputs for each line. The directional shape of the result is robust across a wide range of input assumptions, and the directional shape is what matters.

HOW TO LOCALISE THIS FOR YOUR ORGANISATION

i.

Start with retention

Pull last five years of gross and net dollar retention on accounts acquired through the existing program. The 50%/85% spread is illustrative — your real spread is more useful.

ii.

Layer survival

Cross-reference your founder cohort against the company's actual paying-account survival rates by funding stage. Modest uplifts (1–3 percentage points) are defensible; larger uplifts need explicit mechanism.

iii.

Apply your own ASP

Use your actual average annual spend for a Series A-stage account, segmented by product line. The expansion factor between the two cohorts can be calibrated from NRR on existing high-touch accounts.

iv.

Stress-test the cost line

The $3.5M / $6.0M figures are global-program scale. For a regional or pilot program, the absolute numbers shrink proportionally; the asymmetry of the ratio holds.

05

SECTION

The corporate blind spot

What the dashboard does not measure — and the costs that compound while no one is looking.

§ 05 / 07   APPROX. 5 MIN READ

The hidden costs of the drifted model

The economics of section four describe what a rebuilt program could deliver. This section is about the inverse: what the drifted model is actively costing companies right now, in ways that do not appear on any dashboard the program itself produces.

HIDDEN COST ONE

The strategic intelligence that never reaches leadership

A program that is processing applications and distributing credits is not generating intelligence. It is generating throughput. The signal that a senior executive could once have read off an active program lead's quarterly briefing — which categories are heating up, which architectural patterns are spreading, which competitive threats are forming at the early stage — does not exist in the drifted model. It has been replaced by a pipeline dashboard, which describes the past in commercial categories and tells the executive nothing about the future.

The cost of this is genuinely difficult to size, because it is the cost of not seeing things in time. But it is real. Companies that lose touch with the leading edge of their own ecosystem are routinely surprised by competitive shifts they could have seen coming.

HIDDEN COST TWO

The internal cultural and retention cost

Senior practitioners — engineers, product managers, sales leaders — derive disproportionate satisfaction from helping early-stage founders solve hard problems. It connects them to the mission of their work in a way that quarterly OKRs rarely do. A program that does not offer this opportunity is, in effect, leaving an employee-engagement asset on the table.

The retention math is straightforward: senior technical attrition is expensive, and any intervention that demonstrably reduces it pays back quickly. The drifted model produces no such intervention. The rebuilt model produces it as a structural feature.

RECOGNISE THE SYMPTOMS · IF ANY OF THESE ARE TRUE, THE COSTS ABOVE ARE ACCRUING

- [ ] The startup program has not produced a written intelligence briefing for leadership in the last twelve months.
- [ ] Senior engineers in the company cannot name a single early-stage founder they have spoken to this year.
- [ ] The program lead is measured exclusively on application volume, credits issued, or quarterly pipeline.
- [ ] The accelerator and VC partner list has not been refreshed since the program was reorganised.
- [ ] Internal employees would not describe the program as one of the reasons they stay.

HIDDEN COST THREE

The brand and ecosystem cost

Founder communities are dense and vocal. A founder who has been processed through a portal, ignored for three weeks, and onboarded with a recorded video does not file a complaint. She tells other founders. Over time, those conversations accumulate into a reputation that is difficult to see from inside the company and almost impossible to fix once formed.

The corollary is positive. A company whose program is genuinely well-run earns reputation in the same channels, at the same speed, and with the same compounding. Both effects are invisible to internal dashboards. Both are felt strongly by the founders who eventually become — or do not become — major customers.

> **The drifted program is, in effect, converting a meaningful percentage of the next decade's potential enterprise customer base into competitor advocates, and reporting the process as successful on a quarterly pipeline dashboard.**

Why the program team cannot fix it from below

It is worth stating explicitly why these problems cannot be solved by the people who actually run startup programs day-to-day, however motivated and competent those people are.

A program lead operating under marketing-organisation KPIs cannot redirect the program toward relationship-led engagement, because doing so would tank her quarterly metrics long before the compounding benefits appeared. A lead trying to launch an internal mentorship track cannot do so without cooperation from HR, L&D, and line managers, none of whom have any incentive to prioritise her request. A lead trying to set up a compensated advisory board cannot commit the budget without a sign-off process that will take months.

These are not failures of individual initiative. They are the predictable consequences of the org chart and the incentive structure. The program is sitting where it is sitting, being measured as it is being measured, because the senior executives above it placed it there. Fixing the program requires those same executives to make different choices.

AN INTERLUDE

The dashboard tells you what you measured. Not what mattered.

The next section is about how to measure wh — and rebuild the program around i

30

ALE

06

SECTION

Rebuilding for real engagement

The seven-pillar framework for a program that actually works — and how to start in the first twelve months.

§ 06 / 07 APPROX. 9 MIN READ

FOUNDER NARRATIVE

A different ending.

Maya Chen

`In a parallel world`

In a version of the story where Maya's original cloud vendor had rebuilt their program, her application would never have sat in a portal for eighteen days. A senior program lead would have seen her name come in from a trusted ecosystem partner, reached out within forty-eight hours, and had a thirty-minute conversation about what she was building and where she was stuck.

Within two weeks, Maya would have spent an hour with a senior engineer reviewing her architecture. Within two months, she would have been introduced to another founder in the portfolio solving an adjacent problem.

Three years on, when the Series A round closed, migration to a competitor would not have been the first move she made. She would have stayed. She would have expanded. She would have sent her peer founders over.

Daniel Okafor

`In the real world`

Daniel's experience is not exotic. Priya's company does not have a unique program. They have, simply, a program that is structured differently — run by a founder-minded lead, supported by a compensated advisory board, plugged into a voluntary internal mentorship track, funded across multiple cost centres, and measured over three- to five-year horizons.

None of the components are radical. All of them, taken together, produce a different outcome.

THE SEVEN COMPONENTS PRIYA'S PROGRAM RUNS ON · PREVIEWED HERE, DETAILED IN Section 06

01	02	03	04	05	06	07
Founder-minded program lead	Compensated advisory board	Opt-in internal mentor program	Corporate value map	Cross-functional budget	Long-horizon metrics	Tone & language

FIGURE 2

The seven pillars at a glance

Each pillar addresses a specific failure mode of the drifted model. None works in isolation; the model is the combination.

01 **The founder-minded program lead** The right person, not a marketing hire.	**02** **The compensated advisory board** Quarterly ground-truth from active operators.	**03** **The opt-in internal mentor program** Voluntary engagement. Deliberate matching.
04 **The corporate value map** Lead with what you can genuinely offer. Discounts last.	**05** **Cross-functional budget ownership** Sales, marketing, CSR, and L&D — each with skin in the game.	**06** **Long-horizon, multi-dimensional metrics** Three- to five-year cohort economics. Not quarterly pipeline.

07 **Tone and language** Founders read every word. Write every word as if they will.

~~"for Startups"~~ ⟶ **"Company × Startup"**

06 · REBUILDING FOR REAL ENGAGEMENT

From "for Startups" to "Company × Startup"

The "for Startups" model places the company at the centre and the founder in the position of recipient. The "Company × Startup" model, borrowing from the partnership conventions now ubiquitous in brand and platform culture, places the two parties side by side. Each brings something. Each gains something. Every pillar that follows is, in one way or another, an answer to the question: what would this engagement look like if it were genuinely structured as Company × Startup rather than Company for Startup?

01 The founder-minded program lead

The single most important decision in rebuilding a startup program is the person running it. She does not need to have been a founder herself, although that helps. She does need one specific quality that marketing and sales hiring processes rarely select for: a genuine, sustained interest in founders and early-stage companies, evidenced by her behaviour over years rather than in an interview.

Beyond that orientation, the role demands four capabilities. Highly organised, because every broken promise damages the brand. An excellent brand representative, because her tone in a founder conversation carries the company's tone in that founder's mind for years. Deeply fluent in the internal organisation of her own company, because the core of her job is connecting founders to the right internal people. And strategically literate, because she is the interface between ecosystem intelligence and leadership.

This is not a junior role. Under-levelling it is one of the most common and most damaging mistakes large companies make in their startup functions.

02 The compensated startup advisory board

Every serious program should be advised by three to five founders and ecosystem operators, meeting quarterly, producing a written report, and compensated meaningfully for their time. This is the program's real-time intelligence layer, its reality check, and its primary source of calibration against what is actually happening in ecosystems.

Compensation matters. It ensures that the commitment is taken seriously on both sides, and it removes the awkward power dynamic of asking founders to give substantial time for free to a company that is, in the background, trying to sell to them. The board's output is a quarterly written report covering emerging technology trends, changes in founder behaviour, competitive positioning, and candid feedback on what the program is doing well or badly.

03 The opt-in internal mentor program

Internal employee mentorship of founders is, arguably, the single highest-leverage element of a rebuilt program — and the one most completely absent from the drifted model. The design principles are straightforward.

Employees opt in voluntarily. No mandate, no allocation by manager. People who do not want to mentor do not. People who do, self-select, and that self-selection does almost all the quality control for you. **The commitment is modest and bounded** — five to ten hours per month, time-boxed engagements, mentor controls her own calendar. **The matching is deliberate**: a sales leader matched to a founder struggling with early GTM; a senior engineer matched to a founder evaluating a specific architecture decision; the program lead is the filter, not a self-service portal.

The cultural signalling matters as much as the mechanics. Publicised internally, this program reads as a signal that the company values practitioner expertise — with measurable retention effects among senior staff who might otherwise be looking at adjacent opportunities. A handful of large companies that have run programs like this report mentor-cohort retention rates two to four points higher than matched non-participants — at staff levels where each percentage point is worth hundreds of thousands of dollars in replacement cost.

04 The corporate value map

Before the program can engage founders, it needs to know what it is offering. Most drifted programs have a clear answer to **what do we give founders?** (credits, webinars, a portal login) and no clear answer to **what value can we genuinely provide that founders actually need?** The second answer requires honest internal work that the first sidesteps.

The value-mapping exercise asks: where can this company materially help an early-stage founder? Possible answers include architectural and technical expertise; go-to-market coaching; customer introductions; investor introductions; senior hires; product feedback; operational expertise where the company is genuinely world-class; and yes, discounted access to the company's products. The discount is on the list. It is not at the top.

The advisory board's role here is critical, because internal teams are systematically poor at assessing their own strengths from a founder's point of view. A company may believe its technology is its strongest offering when, in fact, founders would rather access its sales expertise. The board is the external reality check on that gap.

05 Cross-functional budget ownership

The structural fix that makes most of the preceding pillars sustainable is to spread the program's funding across at least four budget lines.

SALES	funds it for long-term customer acquisition and expansion revenue. Justified by five-year CLV economics, not quarterly pipeline.
MARKETING	funds it for brand health, ecosystem positioning, and authentic founder advocacy. Justified by brand-awareness outcomes and founder NPS.
CSR / CORPORATE AFFAIRS	funds it as genuine ecosystem contribution. Justified by ecosystem-contribution metrics and public profile.
L&D	funds it for internal mentorship effects on engagement and retention. Justified by mentor participation, satisfaction, and retention differentials.

The practical effect is governance. With four stakeholders at the table, no single function can compress the program into a funnel that serves only its own metric.

06 Long-horizon, multi-dimensional metrics

A rebuilt program is measured across four dimensions, with the balance heavily weighted toward longer-horizon measures.

COMMERCIAL

Cohort-based, 3–5 year horizon

What percentage of year-one engagement converted into paying accounts by year three; average year-five spend; net dollar retention vs comparable non-program accounts.

RELATIONSHIP

Founder NPS & verbatims

Quarterly NPS across the active cohort, read with particular attention to verbatim responses. A respectable aggregate can mask deeply frustrated text.

INTERNAL

Mentor & retention metrics

Participation rates, mentor satisfaction scores, and retention differentials between participating employees and matched non-participants.

ECOSYSTEM

Positioning & trust signals

Inbound founder interest, partner feedback from accelerators and key VCs, and competitive perception — independently assessed by the advisory board.

07 Tone and language

The tone of a startup engagement program — the specific words used in outreach emails, the framing of the landing page, the warmth or coldness of the onboarding — is not a cosmetic consideration. It is one of the primary signals founders read to determine whether the program sees them as a partner or a lead.

A rebuilt program's voice is personal, specific, and genuinely helpful. It does not lead with credit amounts. It does not say **"our team."** It has names. It responds to individual founders individually. Its forms ask for what is actually needed and no more. Its onboarding introduces a human. Its webinars, when they exist at all, are titled honestly and are not product demos. Its emails are short, direct, and signed by someone whose calendar is linked.

VOICE PRINCIPLES · WHAT FOUNDERS READ ON FIRST CONTACT

DON'T WRITE

- ✕ *"Hi there! Congratulations on your application to our Startup Program."*
- ✕ *"Our team will be in touch within 10 business days."*
- ✕ Generic terms-and-conditions before any human introduction.
- ✕ Subject lines that lead with the dollar value of the credit.
- ✕ Cadenced sequences calibrated to vendor funnel stages.

WRITE INSTEAD

- ✓ *"Maya — saw your application come through from Founders.io. Free for a call Thursday?"*
- ✓ A response time measured in days, signed by a named person.
- ✓ One paragraph of what the program actually offers, in order of usefulness.
- ✓ A subject line that mentions the founder's product by name.
- ✓ Follow-ups paced by what the founder said she needed, not the funnel.

> Founders can tell in the first sentence of the first email.

A FIRST EMAIL, TWO WAYS · THE FOUNDER READS BOTH IN UNDER THREE SECONDS

DRIFTED

Subject: **Welcome to Cloudvendor for Startups — $12,000 in credits inside**

`From: Cloudvendor Startup Team <startups@cloudvendor.com>`

Hi there,

Congratulations on your acceptance to the Cloudvendor for Startups Program. You are now eligible for $12,000 in credits, valid for twelve months. Review the attached

REBUILT

Subject: **Re: Ledgerwise — saw your application from Founders.io**

`From: Priya Anand <priya@cloudvendor.com>`

Maya — thanks for the application. The reconciliation problem you're solving for SMEs is interesting; one of our senior engineers has worked on som

Free for a 30-minute call Thursday?

FIGURE 3

Where to start · the first twelve months

The full benefits compound over two to three years. The first year produces visible change — and the sequence of moves matters.

DAYS 0 – 90

Diagnostic & foundational

- Identify the executive sponsor; commit cross-functional budget in principle.
- Recruit or reappoint the program lead.
- Convene the advisory board.
- Begin the value-mapping exercise.

DAYS 90 – 180

First visible operational changes

- Simplify application; compress time-to-response from weeks to days.
- Launch the internal mentor program with an initial cohort of volunteers.
- Institute founder NPS measurement.
- Produce the first quarterly advisory-board report.

MONTHS 6 – 12

Voice, narrative, & first cohort signal

- Overhaul tone and external communications end-to-end.
- Update internal communication so the program shows up in company narrative.
- Produce the first year-over-year cohort metrics.

None of this is revolutionary. All of it, taken together, produces a program that looks, feels, and performs differently. The founders notice first. The internal employees notice soon after. The commercial outcomes follow.

BY MONTH 12 · WHAT SHOULD BE MEASURABLY DIFFERENT

< 3 days Application-to-first-human response	**100+** Volunteer internal mentors active	**4×** Quarterly intelligence briefings to leadership	**+NPS** First founder-cohort NPS baseline set	**4 / 4** Budget lines (Sales · Marketing · CSR · L&D) co-sponsoring

A SECOND INTERLUDE

None of this is hard.
It is just unfashionable.

40

What follows is a return to first … — and a quiet argument for the patien… them.

07

SECTION

Conclusion

The compounding value of genuine relationships — and why the window is open now.

§ 07 / 07 APPROX. 4 MIN READ

FOUNDER NARRATIVE · FIVE YEARS ON

Where it ends.

Maya Chen

`Series C · 240 employees`

Maya Chen is the CEO of a Series C fintech company with two hundred and forty employees. Her annual software and infrastructure spend is in the low eight figures. She uses, by volume, more of her current cloud provider's services than she ever did of the one she started on.

Her current provider's startup program called her within a week of her original switch and kept in touch for the eighteen months it took to build trust. They remain in touch now.

She does not think about her first cloud provider, except occasionally when asked directly. Her view of them is neutral-to-negative. She does not recommend them to anyone.

Daniel Okafor

`Series C · 310 employees`

Daniel Okafor is the CEO of a Series C productivity company with three hundred and ten employees. His annual software spend is similar to Maya's. He has expanded his relationship with Priya's company into a formal strategic partnership: co-marketing, a shared customer advisory board, a joint reference architecture published to their respective developer communities.

Priya was promoted two years ago; her successor is someone Priya brought in from a founder role, on the advice of the advisory board. Daniel sits on that advisory board now.

He recommends Priya's company constantly, without prompting, because he genuinely believes it is the right choice for any founder solving a problem adjacent to his own.

FIVE YEARS ON · THE COMPOUNDED BALANCE

	Maya	Daniel
Stage today	Series C · 240 employees	Series C · 310 employees
Annual spend with original vendor	Switched away — minimal	Low eight figures · growing
Volume vs early days	Replaced; new vendor now larger	Higher than ever; consolidating
View of original vendor	Neutral-to-negative	Strategic partner
Public advocacy	None · does not recommend	Constant · unprompted
Role in the vendor's program	—	Sits on th

What has been argued

The argument of this paper has been, in essence, simple. Corporate startup engagement programs were designed, originally, to do a specific thing: build genuine, long-term relationships with early-stage founders that would produce, over the following five to ten years, both commercial returns and strategic intelligence for the sponsoring company. They have, over time, drifted into doing something else — distributing discounts through portals to founders who are, in the main, already customers, measured against quarterly pipeline targets that systematically hide the program's real returns and real costs.

The drift was not malicious. It was structural, driven by the gravitational pull of sales and marketing KPIs. It was, and is, reversible. The cost is substantial: lost strategic intelligence, lost employee engagement and retention, and lost long-term customer pipeline measured in the tens of millions of dollars per year for any program running at meaningful scale.

The rebuilt model is not radical. It is, substantially, a rediscovery of what the best early programs already did, combined with a few specific structural innovations — the compensated advisory board, the opt-in internal mentorship, the cross-functional budget, the long-horizon metrics — that address the failure modes of the original model while preserving its strengths. It can be summarised in a single conceptual move: **from "for Startups" to "Company × Startup"** — from a benefaction to a partnership. Everything else follows from that.

Why this matters now

Two developments make the argument more urgent than it would have been a decade ago. The first is the explosion of AI-native startups, a cohort whose technology decisions are being made now, at a speed and scale that has no precedent. The cloud and platform choices a pre-seed AI founder makes in 2026 will, in many cases, shape the infrastructure spend of companies that do not yet have a name but will, by 2030 or 2031, be major enterprise accounts. The window is open. It will close.

The second is the increasing sophistication of founders themselves — better informed, better advised, more deliberate. They are also more vocal, with ecosystem channels that transmit reputational signal at a speed that was simply not possible ten years ago. The cost of running a bad program is no longer invisible. It is publicly priced, in real time, in communities the vendor does not have access to.

WHAT A REBUILT PROGRAM LOOKS LIKE AT YEAR THREE

3–5×	**+2–4 pts**	**< 5%**
Pipeline-to-cost ratio against the discount-only model, measured over a five-year cohort horizon.	Retention lift among senior staff participating in the opt-in mentor program, vs. matched non-participants.	Share of total program spend going to credits. The remainder funds humans, mentorship, and ecosystem participation.
VS. THE DRIFTED BASELINE	CHRO LINE OF SIGHT	CFO LINE OF S

A final thought

There is a version of the argument in this paper that treats it as a matter of commercial optimisation: rebuild the program because the five-year cohort economics are favourable, the CFO will approve the additional spend, the strategic intelligence is a useful byproduct, the employee retention gains are a nice bonus. All of this is true, and all of it is sufficient reason.

But there is a larger frame worth naming before closing. Startup ecosystems are, among other things, one of the most important economic infrastructures of the modern technology industry. They produce the companies that become the customers, the partners, the acquisitions, and occasionally the acquirers, of every major platform incumbent. They are not infinite in their capacity. They are sustained by a fragile web of institutions, intermediaries, mentors, and corporate partners, and they degrade when that web is treated as a funnel rather than a commons.

A large technology company that runs its startup program as a genuine partner to the ecosystems it operates in is not only making a commercial bet. It is participating in the maintenance of the system that makes its own long-term future possible. That participation pays back, in ways that are both measurable and, in the end, more important than measurable. It is worth doing well.

The founders will know. The employees will know. The CFOs will, eventually, know. And the ecosystem will be better for it.

● CONTINUE THE CONVERSATION

Where to go from here

This paper is a compressed version of an argument explored at greater length in a forthcoming book. If any part of it has resonated, sparked disagreement, or raised questions you would like to discuss, the next steps below offer routes back into the conversation.

THE BOOK

We're Doing This Wrong

How to Build Entrepreneur Ecosystems That Actually Work.

AUTHOR	PUBLISHER	PUBLICATION
DC Cahalane	**Cantillion Press**	**November 2026**

Examines how the institutions around founders — corporations, governments, universities, accelerators, and capital providers — can build more productive relationships with the people they exist to support. The corporate–startup engagement argument set out in this white paper is one strand of that larger picture.

SPEAKING & ADVISORY

DC speaks regularly to large technology companies, government bodies, and investor networks on corporate–startup engagement, AI and venture capital, and ecosystem architecture.

`me@dc.ie`

GET IN TOUCH

Reactions — agreement, disagreement, questions, ideas to push back on — are genuinely welcome. The argument is not finished, and contact from readers inside large technology companies has consistently sharpened how it is articulated.

`me@dc.ie` · LinkedIn

● COLOPHON

About the author

DC Cahalane is an Ecosystem Architect — a four-time founder turned VC investor who has built companies in the United States, the United Kingdom, and Ireland.

For over fifteen years he has advised governments and some of the world's largest companies on the entrepreneur economy and how to engage with startups.

He is the founder of Republic of Work, Cork's largest innovation hub, and serves as Venture Capitalist in Residence at University College Cork, Ireland. He also writes the **Founder Therapy** column for the **Business Post**, Ireland's largest business publication.

His forthcoming book, **We're Doing This Wrong: How to Build Entrepreneur Ecosystems That Actually Work**, will be published by Cantillion Press in November 2026.

Get in touch · `me@dc.ie`

25+	4×	3	1
Years working in startup ecosystems	Founder · companies built in US, UK, IE	Continents of ecosystems mapped & advised	Forthcoming book · Cantillion Press, Nov 2026

Lost at Scale

Why your startup engagement isn't really engaging anyone — and how to rebuild it for long-term value.

A white paper by DC Cahalane.

Published June 2026 by Cantillion Press, drawn from the forthcoming book *We're Doing This Wrong: How to Build Entrepreneur Ecosystems That Actually Work*.

Set in Source Serif 4 and Inter Tight, with JetBrains Mono for editorial detail. Designed for both screen and print.