

HOW TO SCALE THE SALES MOTION INTO A **REPEATABLE, MULTI-CHANNEL GROWTH ENGINE**

Growth Acceleration & Unit Economics

WHERE THIS IS USED

- Venture Studio ventures entering scale
- CVC portfolio companies post first tranche
- Corporate venture-client growth programs
- Foundry-as-a-Service engagements

AUDIENCE

- EIR / General Manager
- Growth Operator / Head of Sales
- Executive Sponsor
- Venture Board / IC members

PHASE

Phase Four: Scale & Optimize

Section I — Growth Acceleration & Unit Economics

Follows Phase Three (Build & Launch)

EXECUTIVE SUMMARY

A founder who closed the first ten customers has proven that the product can be sold. It has not proven that the product can be sold repeatably, by someone other than the founder, through more than one channel, at a cost the business can afford. Those are different claims, and the gap between them is where most corporate ventures stall in the first two quarters of scale.

This guide converts the Customer Zero Playbook from Guide F3 and the go-to-market strategy from Guide C3 into an instrumented, multi-channel growth engine. It moves the venture from “we can close deals” to “the system closes deals predictably.” That means a documented motion a cold operator can run, two to four live acquisition channels each with its own measured economics, a Growth Operator who owns the pipeline instead of the EIR, and a funnel instrumented end to end.

The four outputs of this guide are: Growth Playbook v2.0, written for an operator with no prior context; a live channel portfolio with per-channel cost of acquisition and payback; a Growth Operator accountable for the number with the EIR out of the default deal path; and a channel scorecard that ranks channels by scalable, profitable contribution. That scorecard is the primary input to the scaling roadmap in Guide I3 and to the follow-on funding case in Guide J1.



THE CORE PROBLEM

The first ten customers are almost always closed the same way: the EIR knows someone, understands the pain personally, and carries the deal through relationships and intuition. That is exactly right for one through ten. It does not scale, because none of it is written down, none of it is measured, and all of it depends on one person being in every conversation.

The transition from founder-led selling to system-led selling is not a matter of hiring salespeople. It is a matter of turning a tacit motion into an explicit, instrumented, multi-channel engine before the headcount arrives. Ventures that hire first and document later end up with expensive salespeople who have nothing repeatable to run.

The six most consistent failure patterns when scaling the sales motion:

- **Founder dependency.** Every deal still routes through the EIR. The pipeline looks healthy, but it is a single point of failure, and it caps growth at the founder's personal bandwidth.
- **One-channel fragility.** The venture found one channel that worked at the 1–10 stage and treats it as the entire growth plan. When that channel saturates or its cost rises, there is no second engine.
- **Unmeasured channel economics.** Outreach and spend scale, but no one can state the cost of acquisition or payback period by channel. Growth is bought without knowing whether it is profitable.
- **Playbook rot.** The F3 playbook was written for the first ten customers and never updated. The motion that closes customers 11–100 is different, and the document no longer describes what actually works.
- **Hiring ahead of the system.** A sales team is brought on before the motion is documented and the funnel is instrumented. The new hires improvise, produce inconsistent results, and the venture cannot tell whether the problem is the people or the process.
- **Vanity pipeline.** Pipeline value grows, but stage conversion and cycle velocity are not tracked. The reported growth is unqualified, and the forecast the board relies on is fiction.

The real issue:

A venture without a repeatable, instrumented sales motion is not fundable at the next tranche, regardless of how good the product is. The follow-on capital decision in Guide J1 does not turn on revenue alone — it turns on evidence that the engine works without the founder, across more than one channel, at a known and improving cost. If that evidence does not exist, the capital request is a bet on the founder's calendar, not on a scalable business.

3


PREREQUISITES

Must Be Complete Before Starting:

- **Guide F3 completed** — Customer Zero Playbook v1.0 exists, 1–10 customers matching the ICP acquired, and funnel conversion rates documented from first outreach
- **Guide C3 completed** — ICP defined, Buying Center mapped, and the 1-10-100 Channel Framework designed for the priority segment
- **Guide C1 completed** — baseline unit economics modeled (CAC, LTV, contribution margin, payback); this guide feeds the volume data that Guide I2 uses to prove those economics at scale
- **Guide E2 active** — the weekly and monthly operating cadence is running, so growth reviews attach to an existing rhythm rather than a new meeting series

Team Readiness:

- Growth Operator confirmed, named, and dedicated to this motion — not the EIR carrying a second title
- CRM or pipeline system live with defined stages, exit criteria per stage, and a source field on every opportunity for per-channel attribution
- Attribution and analytics instrumentation configured before channel spend is scaled — not retrofitted after the money is spent
- For CVC tracks: the Close Owner has confirmed which growth metrics flow into the PMS and at what cadence



EXPECTED OUTPUT/ SUCCESS CRITERIA

You have completed this guide when the following are true:

- ✓ Growth Playbook v2.0 complete — the motion documented step by step for an operator with no prior context, superseding the F3 v1.0 playbook
- ✓ Two to four acquisition channels live, each with a documented economics profile: cost of acquisition, stage conversion, payback contribution, and estimated scalable ceiling
- ✓ The Growth Operator owns the pipeline and the number; the EIR is out of the default deal path and appears only on named escalations
- ✓ The funnel is instrumented end to end, with per-channel conversion and cycle velocity visible in the CRM and reconciled to the H2 dashboard
- ✓ A channel scorecard ranks every channel by scalable, profitable contribution — the primary input to Guides I3 and J1
- ✓ The founder-dependency test is passed: at least three deals closed in the period with the EIR not in the room



STEP-BY-STEP INSTRUCTIONS

STEP 1

CODIFY THE MOTION INTO GROWTH PLAYBOOK V2.0

Before any channel is scaled or any hire is made, the motion that actually closed the first ten customers must be extracted from the EIR's head and written down as a procedure. Growth Playbook v2.0 is not a refresh of the F3 document — it is a rewrite for a different reader: an operator who was not in any of the first ten deals and has no relationship with the market.

- 1.1 Extract the winning motion from the first ten closes** — For each of the first ten customers, reconstruct the actual sequence: how they were sourced, the first message that got a reply, the number of touches to a meeting, who in the Buying Center moved the deal, the objection that nearly killed it, and the specific commitment that closed it. Patterns across ten deals are the raw material of the playbook.
- 1.2 Rewrite every play for a cold operator** — Rewrite the six F3 playbook sections — ICP, problem and value narrative, outreach and funnel, call structure, offer and contract, objection library — so that a new Growth Operator can execute them without the EIR present. Every instruction that assumes founder context is a defect to be fixed.
- 1.3 Define pipeline stages and hard exit criteria** — Name each pipeline stage and write the objective, evidence-based exit criterion that moves an opportunity to the next stage. A deal advances because a specific commitment was obtained, not because the operator feels optimistic. Ambiguous stages are the source of vanity pipeline.
- 1.4 Version-control the playbook** — Assign a version number, an owner, and a last-updated date. The Growth Operator updates it after every material objection, lost deal, and channel change. A playbook that is not updated monthly is out of date within a quarter.

Playbook Section	F3 v1.0 (Founder-Led)	Growth Playbook v2.0 (System-Led)
ICP and bad-fit	Evidence from 1–10 closes	Firmographic filters an operator can run against a list without judgment
Outreach	EIR's personal network and relationships	Per-channel sequences with timing, templates, and reply hooks
Call structure	Founder tells the origin story	Discovery script, demo flow, and micro-commitment checkpoints anyone can run
Objection library	Held in the EIR's memory	Top 5–7 objections with tested Pause-Probe-Provide-Prove responses
Offer and contract	Bespoke per deal	Standard package: fixed scope, one pricing tier, one contract template

AI PROMPT – Growth Playbook v2.0 Conversion

I am converting a founder-led sales motion into a system-led growth playbook. The venture is [describe]. Here is a summary of how the first ten customers were closed: [paste source, first message, touches, buying-center path, closing commitment for each]. For each of the six playbook sections – ICP and bad-fit, problem and value narrative, outreach and funnel, call structure, offer and contract, objection library – rewrite the section so a new Growth Operator with no relationship to this market can execute it. Flag every instruction that silently depends on founder context, and rewrite it as an explicit, repeatable step. Then propose pipeline stages with an objective exit criterion for each.

STEP 2 INSTRUMENT THE FUNNEL AND DEFINE CHANNEL ATTRIBUTION

A motion you cannot measure per channel cannot be scaled deliberately. Before adding channels, extend the F3 funnel so every opportunity carries its source and every stage transition is tracked. Attribution decided after the spend is guesswork; attribution decided before the spend is a control system.

- 2.1 **Extend the F3 funnel with per-channel tracking** — Keep the F3 funnel stages — outreach, reply, meeting, pilot commit, paid conversion — and add a mandatory source field so each metric can be sliced by channel. The same funnel, resolved by channel, is what turns “we are growing” into “this channel converts at this rate for this cost.”
- 2.2 **Choose an attribution model and write it down** — For early scale, first-touch attribution by channel is usually sufficient and unambiguous. Whatever model is chosen, define it once, record it in the metric definition register, and apply it consistently. Changing attribution silently makes every trend meaningless.

2.3

Reconcile the funnel to the H2 dashboard — The commercial metrics in the funnel must reconcile to the Commercial Progress category of the H2 dashboard. One number, one definition, one source — the board pack and the growth review cannot report different pipeline figures.

Funnel Stage	Metric	Tracked By Channel?	What A Drop Signals
Outreach / reach	Prospects contacted or reached	Yes	Sourcing volume or channel capacity is the constraint
Reply / lead	Replies or inbound leads	Yes	Message, creative, or targeting is off for that channel
Meeting	Qualified meetings held	Yes	The hook converts interest but not attention
Pilot / opportunity	Pilots or opportunities created	Yes	Wrong Buying Center contact or wrong offer for the stage
Paid conversion	Signed paying customers	Yes	Commercial path or pricing friction in that channel

AI PROMPT — Funnel Instrumentation and Attribution Design

I am instrumenting a multi-channel sales funnel for a scaling venture. Current single-channel funnel metrics from Guide F3: [paste reply rate, meeting rate, pilot commit rate, paid conversion]. Channels I intend to run: [list]. Design the funnel so every stage can be resolved by channel, recommend a first attribution model with its trade-offs, and list the exact CRM fields required on every opportunity to make per-channel conversion and cycle velocity reportable without manual reconciliation.

STEP 3

DIVERSIFY CHANNELS THROUGH CONTROLLED EXPERIMENTS

A single working channel is a fragile growth plan. The objective of this step is not to run every channel at once — it is to test two or three additional channels as bounded experiments, each with a budget cap, a hypothesis, and an explicit kill criterion, so that a portfolio of proven channels replaces a single point of failure.

3.1

Select candidate channels from the C3 framework — Return to the 1-10-100 Channel Framework in Guide C3. The channels that fit customers 100+ are the candidates here. Prioritize channels whose economics could plausibly beat the current channel, not channels that are simply available.

3.2

Design each channel test as a bounded experiment — Each test states a hypothesis, a fixed budget or time cap, a target cost of acquisition, and a kill criterion. A channel that has not beaten its kill criterion within the cap is stopped — not extended in hope. Controlled, reversible bets are the discipline that keeps scale spend evidence-based.

3.3

Run one channel experiment at a time to isolate signal — Running three new channels simultaneously makes attribution noisy and wastes budget. Sequence the experiments so each channel's economics can be read cleanly before the next begins.

Channel	Best Fit	Typical Motion	Watch-Out
Outbound / SDR	Defined ICP, mid-market and up	Sequenced email and social plus qualified meetings	Cost rises fast; needs constant list and message hygiene
Inbound / content	Broad awareness, longer cycles	Content, search, and demand capture	Slow to prove; attribution is hardest here
Partner / corporate distribution	Access to the Champion's network	Co-sell, referrals, procurement fast-lane from Guide E3	Depends on partner incentives; can stall without a named owner
Marketplace / self-serve	Simple product, low-touch buyer	Listing plus product-led onboarding	Only works if the product sells itself
Referral	Delighted early customers	Structured asks after a proven outcome	Needs a real outcome first; not a cold-start channel

Note: For CVC portfolio companies, the corporate parent's distribution is frequently the highest-leverage channel and the reason the investment thesis exists. Activate it deliberately through the Champion relationships and the E3 access, and measure it with the same rigor as any paid channel — a distribution channel that is assumed rather than measured is the one most likely to underdeliver.

AI PROMPT — Channel Experiment Design

I am designing controlled experiments to diversify a scaling venture's acquisition channels. Current working channel and its economics: [describe]. Candidate channels from the C3 framework: [list]. For each candidate, write a one-line hypothesis, a recommended budget or time cap, a target cost of acquisition to beat, and an explicit kill criterion. Then recommend the sequence in which to run the experiments so each channel's economics can be read without contamination from the others.

STEP 4

MEASURE CHANNEL ECONOMICS AND BUILD THE CHANNEL SCORECARD

Every live channel must be reduced to four numbers: fully loaded cost of acquisition, stage conversion, payback contribution, and estimated scalable ceiling. Those numbers, side by side, are the channel scorecard — the single artifact that tells the venture where the next dollar of scale spend should go and, just as importantly, where it should not.

- 4.1 **Calculate fully loaded cost of acquisition per channel** — Include media or list cost, the loaded cost of the operator time spent, and tooling. A channel that looks cheap on media alone is often expensive once operator hours are counted. Use the same cost definitions as Guide C1 so the numbers reconcile to the unit-economics work in Guide I2.
- 4.2 **Estimate the scalable ceiling of each channel** — A channel with excellent economics but a low volume ceiling is a supplement, not an engine. Estimate how many customers per period the channel can realistically produce before its cost of acquisition degrades, and record it — ceiling is what separates a growth engine from a nice-to-have.
- 4.3 **Assign a verdict to each channel** — Every channel gets one of three verdicts: scale, hold, or cut. Scale channels beat their target economics and have headroom. Hold channels work but have limited ceiling or unproven economics. Cut channels have failed their kill criterion. No channel stays live without a verdict.

Channel	Loaded CAC	Paid Conversion	Payback	Scalable Ceiling	Verdict
[Channel A]	[value]	[rate]	[months]	[per quarter]	Scale
[Channel B]	[value]	[rate]	[months]	[per quarter]	Hold
[Channel C]	[value]	[rate]	[months]	[per quarter]	Cut

Note: The scorecard is where Guide I1 hands off to Guide I2. I1 measures channel economics as observed; I2 stress-tests whether those economics hold at volume and clear the scale-worthy bar. Do not declare a channel proven on a small sample — flag it for the I2 cohort analysis.

AI PROMPT — Channel Scorecard Construction

I am building a channel scorecard for a scaling venture. For each live channel I will provide spend, operator hours, opportunities, and closed customers over the measurement period: [paste]. Using cost definitions consistent with Guide C1, calculate fully loaded cost of acquisition, paid conversion rate, and payback period per channel. Estimate a plausible scalable ceiling for each and flag the assumptions behind it. Recommend a scale, hold, or cut verdict per channel with the reasoning, and identify which channels have too small a sample to trust without the I2 cohort analysis.



STEP 5

REMOVE FOUNDER DEPENDENCY AND TRANSFER
OWNERSHIP TO THE GROWTH OPERATOR

A motion that only the founder can run is not a growth engine. This step transfers the motion to the Growth Operator through a deliberate shadow-to-solo progression and proves the transfer with an explicit test, so the EIR's calendar stops being the ceiling on growth.

5.1

Shadow, then co-sell, then solo — The Growth Operator first shadows the EIR on live deals, then co-sells with the EIR present but silent, then runs deals solo with the EIR available only for named escalations. Each stage is complete when the operator can run it without prompting.

5.2

Define the escalation boundary — Write down exactly which situations justify pulling in the EIR — for example, an enterprise deal above a threshold, a strategic account, or a novel objection. Everything else is the operator's to run. An undefined boundary quietly re-centers every deal on the founder.

5.3

Pass the founder-dependency test — The transfer is proven when at least three deals in the period close with the EIR not in the room and not on the calls. Until that is true, the venture has a talented founder, not a growth engine, and Guide J1 will treat the pipeline as founder-dependent risk.

AI PROMPT — Founder-Dependency Transition Plan

I am transferring a founder-led sales motion to a dedicated Growth Operator. Current state: [describe how the EIR is involved in deals today]. Design a shadow-to-solo transition plan with clear completion criteria for each stage, a written escalation boundary that names exactly when the EIR should re-enter a deal, and a founder-dependency test that proves the transfer is real. Identify the two or three deal types most likely to keep re-centering on the founder and how to hand each one off cleanly.

STEP 6

SET GROWTH TARGETS AND THE OPERATING CADENCE FOR SCALE

Scale without a cadence produces effort, not compounding. Attach growth to the existing E2 operating rhythm with a weekly growth review, a forward-looking pipeline coverage discipline, and targets the Growth Operator owns and the board can hold them to.

6.1

Run a weekly growth review inside the E2 cadence — Add a standing weekly review of the funnel by channel, the channel scorecard, and the pipeline coverage ratio. Keep it inside the E2 stand-up and milestone rhythm rather than creating a parallel meeting series that competes for attention.

- 6.2 Enforce a pipeline coverage ratio** — Set and monitor the ratio of qualified pipeline to the period target — commonly three to four times target for early-stage motions. Coverage below the ratio is an early warning that this quarter's number is at risk, visible weeks before it shows up in closed revenue.
- 6.3 Set targets the operator owns** — Define per-channel and total targets for the next two to three periods, owned by the Growth Operator, with the assumptions stated. These become the base case for the I3 scaling roadmap and the growth line item in the J1 capital request.

Cadence	Focus	Owner	Output
Weekly	Funnel by channel, coverage ratio, stuck deals	Growth Operator	Actions and reforecast
Monthly	Channel scorecard review, verdicts, CAC trend	Growth Operator / EIR	Scale / hold / cut decisions in the milestone review
Quarterly	Growth against target, evidence for capital	EIR	Input to Venture Board pack and Guide J1

STEP 7

PACKAGE THE GROWTH EVIDENCE FOR FOLLOW-ON
CAPITAL AND THE BOARD

The final step turns the engine into evidence. The follow-on funding case in Guide J1 asks one question of the growth motion: does it work repeatably, without the founder, at an improving cost, with headroom to scale? Assemble the answer as a short, defensible package rather than a narrative.

- 7.1 Assemble the growth evidence package** — The package is four things: Growth Playbook v2.0, the instrumented funnel by channel, the channel scorecard with verdicts, and the founder-dependency test result. Together they demonstrate a system, not a founder — which is exactly what the next tranche is underwriting.
- 7.2 Fit the evidence to the H2 dashboard and H3 board pack** — Feed the channel scorecard and funnel into the Commercial Progress category of the H2 dashboard so the board sees the same numbers the growth review sees. The quarterly summary drops directly into the H3 board pack.
- 7.3 Apply the dual-track split for CVC portfolio companies** — For a Venture Studio build-and-launch venture, the Growth Operator sits inside the venture team and the evidence flows through the EIR to the Venture Board. For a CVC portfolio company operating independently, the Growth Operator sits inside the portfolio company; the Close Owner reconciles the growth metrics into the PMS at the agreed cadence and carries them to the Investment Committee. In both cases the evidence bar is identical — only the reporting path differs, and the split should already be recorded in the H1 Expectations Alignment Record.

AI PROMPT – Growth Evidence Package for Follow-On Capital

I am assembling the growth evidence for a follow-on capital request under Guide J1. Inputs: Growth Playbook v2.0 status, the instrumented funnel by channel, the channel scorecard with verdicts, and the founder-dependency test result: [paste each]. Draft a one-page growth evidence summary that answers, with numbers, whether the motion is repeatable, founder-independent, improving in cost, and has scalable headroom. Flag the single weakest point in the evidence that the Investment Committee is most likely to challenge, and what would strengthen it.



TROUBLESHOOTING

Symptom	Likely Cause	Fix
Pipeline value is growing but closed revenue is flat	Vanity pipeline: stages advance on optimism, not on evidence-based exit criteria	Enforce the Step 1.3 exit criteria. Re-stage the entire pipeline against objective commitments. Discount or remove any opportunity that cannot show the evidence its stage requires.
One channel produces almost all deals and the others are quiet	Channel experiments were never run as bounded tests, or were abandoned before their cap	Return to Step 3. Run the next candidate channel as a single, isolated experiment with a budget cap and kill criterion. Do not treat the working channel as the growth plan.
Cost of acquisition looks cheap but the business is not profitable	CAC is measured on media spend only and excludes loaded operator time	Recalculate fully loaded CAC per Step 4.1 using the C1 cost definitions. Include operator hours and tooling. Reconcile to the Guide I2 unit-economics work before trusting the number.
Every deal still needs the EIR to close	The motion was never transferred; the escalation boundary is undefined	Run the Step 5 shadow-to-solo progression and write the escalation boundary. Do not report the motion as scalable until the founder-dependency test passes.
New Growth Operator's results are inconsistent and hard to diagnose	Hiring happened ahead of a documented, instrumented system	Complete Growth Playbook v2.0 and the funnel instrumentation before expecting consistent output. With the system in place, inconsistency becomes a coachable, measurable gap rather than a mystery.
The board pack and the growth review report different pipeline numbers	The funnel is not reconciled to the H2 dashboard; two sources exist	Apply Step 2.3. Make the CRM the single source, feed it to the H2 Commercial Progress category, and retire any parallel spreadsheet.
Corporate distribution channel was assumed but produced little	The channel was treated as automatic rather than measured and owned	Assign a named owner, activate the E3 access and Champion relationships explicitly, and measure it on the scorecard like any paid channel. If it still underperforms, give it a hold or cut verdict.

VALIDATION STEPS

Confirm each of the following before declaring the sales motion scaled and ready for the follow-on capital case:

Growth Playbook v2.0 complete, version-controlled, and executable by an operator with no prior context



Pipeline stages have objective, evidence-based exit criteria and the pipeline has been re-staged against them



The funnel is instrumented end to end, resolvable by channel, and reconciled to the H2 dashboard



Two to four channels are live, each tested as a bounded experiment with a documented result



Fully loaded cost of acquisition, payback, conversion, and scalable ceiling are recorded for every live channel



The channel scorecard assigns a scale, hold, or cut verdict to every channel




The escalation boundary is written down and the founder-dependency test has passed with at least three founder-absent closes



A weekly growth review runs inside the E2 cadence and a pipeline coverage ratio is enforced



The growth evidence package is assembled and fits the H2 dashboard and H3 board pack; for CVC tracks, the PMS reporting path is confirmed




NEXT STEPS

With a repeatable, instrumented, multi-channel motion in place, proceed to Guide I2 to prove the unit economics at volume. The channel scorecard from this guide is the input: I2 stress-tests whether the observed economics hold across cohorts and clear the scale-worthy bar before capital is requested.

Feed the channel scorecard and growth targets into Guide I3, which sequences growth investment and operational capacity against the evidence and produces the scaling roadmap.

The growth evidence package assembled in Step 7 becomes a primary input to Guide J1, the follow-on funding and capital deployment case. Keep it current in the H2 dashboard so the capital request is always one export away, not a quarter-end scramble.



MASTER CHECKLIST

A. PLAYBOOK CODIFICATION

- ☐ Winning motion extracted from all first-ten closes (source, first message, touches, buying-center path, closing commitment)
- ☐ All six playbook sections rewritten for a cold operator; founder-context defects removed
- ☐ Pipeline stages defined with objective, evidence-based exit criteria
- ☐ Growth Playbook v2.0 version-numbered, owned, and dated

B. FUNNEL INSTRUMENTATION

- ☐ F3 funnel extended with a mandatory source field on every opportunity
- ☐ Attribution model chosen, documented, and recorded in the metric definition register
- ☐ Per-channel conversion and cycle velocity reportable without manual reconciliation
- ☐ Funnel reconciled to the H2 dashboard Commercial Progress category

C. CHANNEL DIVERSIFICATION

- ☐ Candidate channels selected from the C3 1-10-100 framework
- ☐ Each channel test designed with hypothesis, budget or time cap, target CAC, and kill criterion
- ☐ Experiments sequenced and run one at a time to isolate signal
- ☐ For CVC tracks: corporate distribution activated via E3 and Champion relationships and measured, not assumed

D. CHANNEL ECONOMICS

- ☐ Fully loaded CAC calculated per channel using C1 cost definitions
- ☐ Scalable ceiling estimated and assumptions recorded for each channel
- ☐ Scale / hold / cut verdict assigned to every live channel
- ☐ Small-sample channels flagged for the I2 cohort analysis

E. FOUNDER INDEPENDENCE

- ☐ Shadow-to-solo transition completed with stage-by-stage completion criteria
- ☐ Escalation boundary written down and agreed
- ☐ Founder-dependency test passed: at least three deals closed with the EIR not in the room

F. CADENCE AND TARGETS

- ☐ Weekly growth review running inside the E2 cadence
- ☐ Pipeline coverage ratio set and monitored
- ☐ Per-channel and total targets set for the next two to three periods, owned by the Growth Operator

G. EVIDENCE PACKAGE

- ☐ Growth evidence package assembled: playbook, funnel, scorecard, founder-dependency result
- ☐ Package fits the H2 dashboard and H3 board pack
- ☐ For CVC tracks: PMS reporting path and cadence confirmed with the Close Owner

H. FINAL CHECK

- ☐ The motion runs without the EIR in the default deal path
- ☐ Growth is bought at a known, improving cost across more than one channel
- ☐ The evidence would survive Investment Committee scrutiny at Guide J1

If growth still depends on the founder's calendar
→ **The engine is not yet built. Return to Step 5.**

