GUIDE I3

TURN8

HOW TO BUILD AN **EVIDENCE-BASED SCALING ROADMAP** AND **CAPACITY PLAN**

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
- Finance Lead / Growth Operator
- Executive Sponsor
- Venture Board / IC members

PHASE

Phase Four: Scale & Optimize
Section I — Growth Acceleration & Unit Economics
Follows Phase Three (Build & Launch)

1

EXECUTIVE SUMMARY

By this point the venture has proven two things: that the sales motion is repeatable across channels without the founder (Guide I1) and that the unit economics are scale-worthy in the channels and segments worth scaling (Guide I2). What it does not yet have is a plan for the order in which to spend the next dollars and add the next people so that growth compounds instead of breaking the business.

This guide produces the Scaling Roadmap: an evidence-based, sequenced plan that ties each growth investment to the operational capacity required to support it, identifies the constraint that will bind first, and uses predictive analytics to prioritize where the next dollar of scale spend goes. It replaces a roadmap built on ambition with one built on the I1 and I2 evidence and the C2 financials.

The four outputs of this guide are: a named, tested scaling constraint; a capacity plan whose hires and investments are triggered by leading indicators rather than the calendar; a quarter-by-quarter sequenced roadmap with a decision gate at each stage; and a set of leading indicators reconciled to runway. The Scaling Roadmap is the primary input to the follow-on funding case in Guide J1 — the artifact the Investment Committee funds.

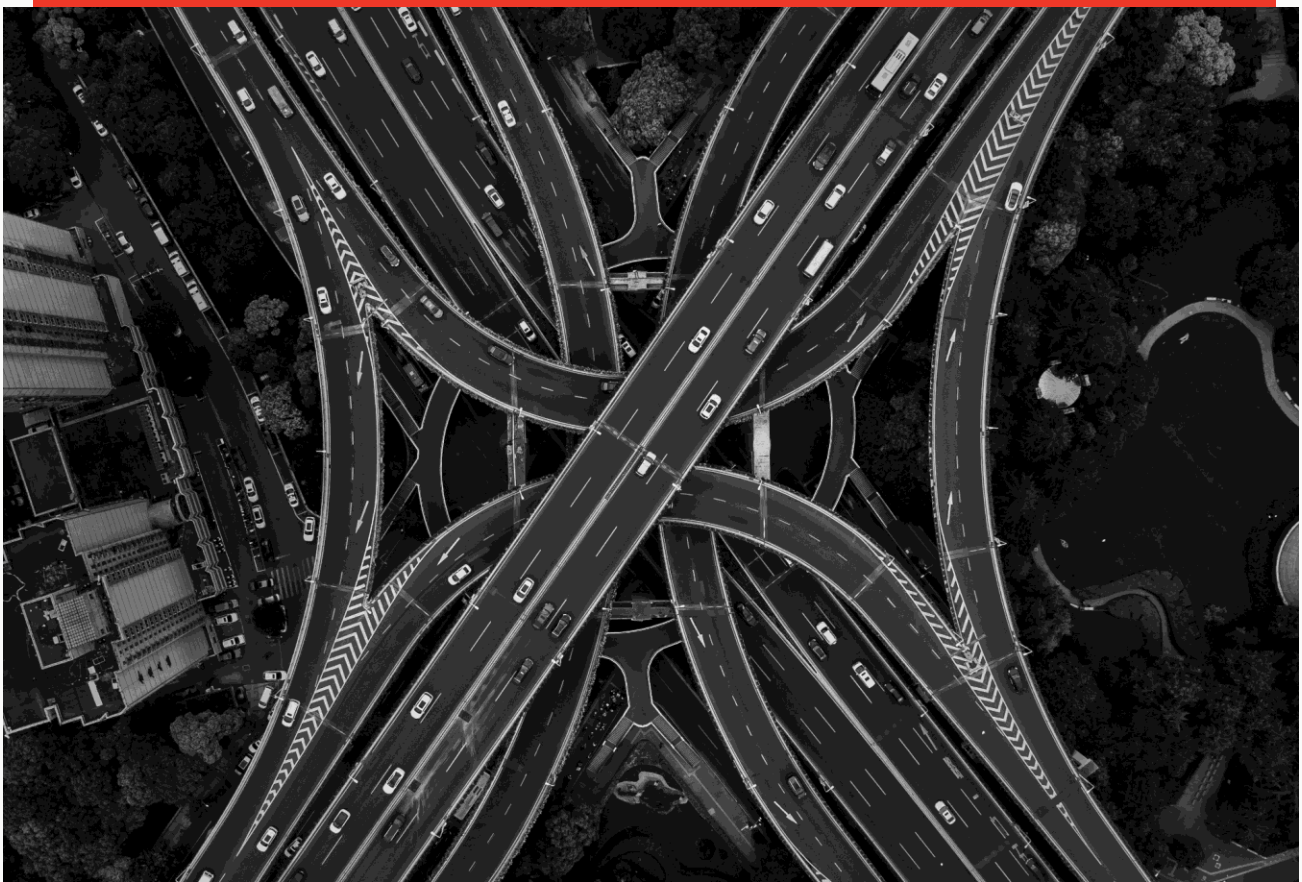

2

THE CORE PROBLEM

Ventures rarely fail at scale because they lacked ambition. They fail because they scaled the wrong thing in the wrong order — pouring money into demand while delivery quietly fell apart, or hiring against a plan that no evidence supported, or committing the whole tranche to a single bet. A roadmap is a sequence of reversible decisions under a runway constraint, and treating it as a straight line of ambition is what breaks businesses that had every reason to succeed.

The discipline this guide imposes is simple to state and hard to hold: every item on the roadmap must earn its place with evidence, be triggered by a leading indicator, and be reconciled to runway before it is committed.

The six most consistent failure patterns in scale-phase roadmaps:

- **Roadmap by ambition.** Targets are set top-down from what the venture wishes were true, disconnected from the I1 channel evidence and the I2 economics.
- **Growth outruns capacity.** Sales scale faster than delivery, support, or operations, so newly acquired customers churn and the growth destroys the retention the model depended on.
- **Hiring by gut.** Headcount is added on instinct rather than against the binding constraint, producing cost without relieving the thing that was actually limiting growth.
- **One big bet.** The entire scale budget is committed to a single channel or segment instead of a sequence of staged, reversible bets that can be stopped when the evidence turns.
- **No leading indicators.** The roadmap is tracked on lagging revenue, so a problem is visible only after a quarter has already been lost.
- **Runway blind.** The plan is never reconciled to burn and runway, and the venture scales confidently into a wall it could have seen from the start.

The real issue:

The follow-on capital decision in Guide J1 funds a roadmap, not a hope. The Investment Committee deploys capital against a sequenced plan with named constraints, capacity triggers, and predicted outcomes it can hold the venture to — or it does not deploy. A scaling plan that cannot survive that scrutiny is not a plan for growth; it is a request to fund uncertainty.

PREREQUISITES

Must Be Complete Before Starting:

- **Guide I1 completed** — the channel scorecard exists with scale, hold, and cut verdicts and a repeatable, founder-independent motion
- **Guide I2 completed** — unit economics proven scale-worthy, with the profitable-growth path classifying every segment and channel
- **Guide C2 updated** — the financial model, burn rate, and runway are current and reconcile to actuals
- **Guides E1 and E2 in place** — the venture team is deployed and the operating cadence is live, so the roadmap attaches to an existing rhythm and org

Team Readiness:

- EIR or General Manager available to own the roadmap and its gates
- Finance Lead available to reconcile the plan to burn and runway
- Historical funnel, delivery, and operations data accessible to seed the predictive baselines
- For CVC tracks: the Close Owner has confirmed how roadmap milestones and leading indicators feed the PMS and the J3 portfolio-health monitoring

4

EXPECTED OUTPUT/ SUCCESS CRITERIA

You have completed this guide when the following are true:

- The binding scaling constraint is named and tested — the venture knows what will break first
- A capacity plan for hiring, infrastructure, and operations exists, with each investment tied to a leading-indicator trigger and a named owner
- A predictive prioritization ranks scale investments by expected return and payback across a base, upside, and downside scenario
- A quarter-by-quarter Scaling Roadmap sequences growth bets against capacity, each with a decision gate
- Leading indicators are defined and reconciled to the C2 runway, with early-warning thresholds
- The Scaling Roadmap is packaged as the primary input to Guide J1; for CVC tracks, the PMS and J3 monitoring linkage is confirmed

STEP-BY-STEP INSTRUCTIONS

STEP 1 ESTABLISH THE EVIDENCE BASELINE

The roadmap is only as good as the evidence under it. Before sequencing anything, assemble the three evidence sources into one baseline so that every subsequent decision refers to the same facts: what grows, what pays, and what the runway allows.

1.1 **Pull the channel scorecard from Guide I1** — Bring in the scale, hold, and cut verdicts and the scalable ceiling per channel. These bound what the roadmap can realistically ask any channel to deliver.

1.2 **Pull the profitable-growth path from Guide I2** — Bring in the scale, fix, and cut classification by segment and channel. The roadmap sequences the scale cases, schedules the fixes, and excludes the cuts.

1.3 **Pull burn and runway from Guide C2** — Bring in the current burn rate and runway. Runway is the hard constraint the entire roadmap is sequenced within; a plan that ignores it is not a plan.

> **AI PROMPT — Scaling Evidence Baseline**
>
> ```
> I am assembling the evidence baseline for a scaling roadmap. From
> Guide I1, the channel scorecard: [paste verdicts and scalable
> ceilings]. From Guide I2, the profitable-growth path: [paste
> scale/fix/cut by segment and channel]. From Guide C2, current burn
> and runway: [paste]. Synthesize these into a one-page baseline that
> states, in numbers, what the venture can profitably scale, how far
> each channel can go, and how much runway is available to do it. Flag
> any contradiction between the three sources that must be resolved
> before roadmapping.
> ```

STEP 2 IDENTIFY THE BINDING SCALING CONSTRAINT

Something will break first as the venture grows — demand generation, delivery capacity, or capital. Scaling everything at once wastes the scarce resource on the things that were not limiting. This step diagnoses the single binding constraint so the roadmap can relieve it before pushing harder on everything else.

Candidate Constraint	Symptom If Binding	How To Test	If It Is Binding
Demand	Pipeline coverage below target; channels near ceiling	Check I1 coverage ratio and channel ceilings	Invest in channel expansion before delivery
Delivery / product	Onboarding backlog; rising churn in new cohorts	Check I2 cohort retention for newest cohorts	Add delivery capacity before adding demand
Operations / support	Response times slipping; support cost per customer rising	Check cost-to-serve trend from I2	Invest in ops and tooling before scaling volume
Capital	Payback longer than runway allows for the plan	Reconcile roadmap spend to C2 runway	Sequence smaller staged bets; this is the J1 case

2.1 **Test each candidate against the evidence, not intuition** — Run each candidate constraint against the baseline data from Step 1. The constraint is the one whose symptom is already visible or arrives soonest at the planned growth rate.

2.2 **Name one binding constraint at a time** — There is usually one constraint that binds first. Relieving it will surface the next one — that is expected and healthy. A roadmap that tries to relieve all constraints simultaneously relieves none decisively.

STEP 3 MODEL THE NEXT DOLLAR WITH PREDICTIVE PRIORITIZATION

With the constraint named, the question becomes where the next unit of scale spend produces the most profitable growth. This is where predictive analytics earns its place: using the venture's own funnel, retention, and cost history to forecast the expected return of each candidate investment, rather than allocating by advocacy or recency.

3.1 **Build predictive baselines from historical data** — Use the historical funnel, cohort retention, and cost-to-serve data to fit simple, explainable forecasts: expected conversion and payback per channel, expected retention per segment, and expected cost-to-serve as volume grows. Favor models the team can interpret and defend to the Investment Committee over opaque ones.

3.2 **Rank candidate investments by expected return and payback** — For each candidate scale investment, project the incremental customers, contribution margin, and payback using the baselines. Rank them. The ranking, not a hunch, decides where the next dollar goes.

3.3 **Model base, upside, and downside scenarios** — Run the ranked plan under three scenarios, varying the two inputs the model is most sensitive to. The downside scenario is the one that must still fit within runway — if it does not, the plan is too aggressive regardless of how good the base case looks.

> **AI PROMPT — Predictive Prioritization of Scale Spend**
> `I am prioritizing scale investments for a venture using predictive analytics. The binding constraint is [state it]. Historical data: funnel conversion and CAC by channel, cohort retention by segment, and cost-to-serve versus volume: [paste]. Build explainable forecasts for conversion, payback, retention, and cost-to-serve, then project incremental customers, contribution margin, and payback for each candidate investment: [list]. Rank the candidates by expected return within payback and runway limits, and model base, upside, and downside scenarios — flagging whether the downside still fits our C2 runway of [runway].`

STEP 4 BUILD THE CAPACITY PLAN WITH TRIGGERS

Capacity is added to relieve the constraint, in step with growth, and only when a leading indicator says it is needed. Trigger-based capacity replaces calendar-based hiring: you commit the hire or the investment when the metric crosses the threshold, not when the quarter starts. This is what keeps the venture from hiring ahead of evidence or scaling into a wall behind it.

Capacity Area	Trigger Metric & Threshold	Investment	Owner
Growth / sales	Pipeline coverage falls below target for 2 weeks	Add Growth Operator capacity or channel spend	Growth Operator
Delivery / onboarding	Onboarding time or new-cohort churn crosses threshold	Add delivery or customer-success capacity	EIR / GM
Operations / support	Cost to serve or response time crosses threshold	Add tooling, automation, or support capacity	Operations owner
Infrastructure	Usage or reliability metric crosses threshold	Add infrastructure or platform capacity	Technical lead

4.1 **Tie every hire and investment to a leading indicator** — For each capacity item, write the specific metric and threshold that triggers it and the owner accountable for acting. A capacity plan without triggers is just a hiring wish list that spends runway on the calendar's schedule.

4.2 **Sequence capacity ahead of the demand it must serve** — Delivery and support capacity are added slightly ahead of the demand they will serve, because capacity that arrives after the customers churns them. Demand capacity is added only against the channels the scorecard says can absorb it.

Note: The Phase Four AI-enhancement opportunity applies directly here: the predictive baselines from Step 3 can drive the leading-indicator thresholds in this capacity plan and, later, feed the automated portfolio-health monitoring in Guide J3. Build the indicators once so they serve both the roadmap and the monitoring.

STEP 5 SEQUENCE THE ROADMAP

The roadmap sequences the ranked growth bets and the triggered capacity into a quarter-by-quarter plan, with a decision gate at each stage. Each stage is a staged, reversible bet: it proceeds to the next only if the gate is cleared, which is what makes the plan fundable rather than a single all-or-nothing commitment.

5.1 **Lay out the bets and capacity by quarter** — For each of the next two to four quarters, state the growth bet, the capacity added to support it, the decision gate, and the expected outcome. Front-load the highest-ranked, constraint-relieving moves.

5.2 **Define a gate at every stage** — Each stage has an evidence-based gate — a specific result that must be true to release the next stage's spend. A gate that is not cleared pauses or redirects the roadmap; it does not get waved through on momentum. This mirrors the Go/Hold/Stop discipline the venture already knows.

Quarter	Growth Bet	Capacity Added	Decision Gate	Expected Outcome
Q1	Scale the top-ranked channel	Growth + onboarding capacity	CAC and payback hold at 2× volume	Proven scalable channel
Q2	Add second scale channel; fix one segment	Delivery + support capacity	Second channel clears the bar	Two proven engines
Q3	Expand proven channels	Ops + infrastructure capacity	Retention stable at higher volume	Compounding, stable growth
Q4	Press advantage within runway	Trigger-based only	Downside scenario still fits runway	Roadmap ready to refinance at J1

AI PROMPT — Roadmap Sequencing

```
I am sequencing a scaling roadmap for a venture. The binding
constraint is [state it], the ranked investments from the predictive
model are [paste], the trigger-based capacity plan is [paste], and
the C2 runway is [runway]. Sequence these into a quarter-by-quarter
roadmap for the next [N] quarters. For each quarter, state the growth
bet, the capacity added, an evidence-based decision gate that must
clear to release the next stage, and the expected outcome. Ensure the
cumulative spend never breaches runway in the downside scenario, and
mark the quarter at which the roadmap is ready to support a J1
capital request.
```

STEP 6 SET LEADING INDICATORS AND RECONCILE TO RUNWAY

A roadmap tracked on lagging revenue reveals problems a quarter too late. Define the leading indicators that will show the plan working or failing weeks ahead of the revenue, wire them to the H2 dashboard, and reconcile the whole plan to burn and runway so the venture never scales past its cash.

6.1 **Define leading indicators per bet and per constraint** — For each growth bet, name the early metric that moves first — pipeline coverage, reply rate, onboarding time, new-cohort retention. These are the tripwires; the capacity triggers in Step 4 are wired to them.

6.2 **Reconcile the roadmap to C2 burn and runway** — Roll the roadmap spend and expected revenue into the C2 model and confirm the runway holds in the downside scenario. If it does not, cut or defer the lowest-ranked bets until it does. This reconciliation is what turns the roadmap into a credible capital ask.

6.3 **Wire the indicators to the H2 dashboard** — Put the leading indicators and gate status on the H2 dashboard so the Venture Board and, for CVC tracks, the Investment Committee see the roadmap's health live, not at quarter-end. This is the same feed that Guide J3 will later use for automated portfolio-health monitoring.

STEP 7 PACKAGE THE SCALING ROADMAP FOR FOLLOW-ON CAPITAL

The Scaling Roadmap is the artifact Guide J1 funds. The final step packages it so the Investment Committee can see exactly what the capital buys, in what order, against which evidence, with which gates, and within which runway.

7.1 **Assemble the scaling roadmap package** — The package is the evidence baseline, the named constraint, the predictive prioritization with scenarios, the trigger-based capacity plan, the sequenced roadmap with gates, and the runway reconciliation. It reads as a plan a committee can hold the venture to, not a narrative of intent.

7.2 **State what each tranche of capital unlocks** — Map the capital request to the roadmap gates: this tranche funds Q1–Q2 to clear these gates; the next tranche is released on that evidence. Staged capital against staged gates is the structure Guide J1 is built to evaluate.

7.3 **Apply the dual-track split for CVC portfolio companies** — For a Venture Studio build-and-launch venture, the EIR owns the roadmap and it flows through the Venture Board; capacity is added inside the venture team. For an independent CVC portfolio company, the portfolio company's management owns the operating roadmap while the Close Owner tracks its milestones and leading indicators in the PMS, which feeds the J3 portfolio-health monitoring and the Investment Committee's view. The evidence discipline and gates are identical across both tracks; only the ownership and reporting path differ, and the split should already be recorded in the H1 Expectations Alignment Record.

AI PROMPT — Scaling Roadmap Package for Follow-On Capital

```
I am packaging a scaling roadmap as the primary input to a follow-on
capital request under Guide J1. Inputs: evidence baseline, named
constraint, predictive prioritization with base/upside/downside
scenarios, trigger-based capacity plan, sequenced roadmap with gates,
and C2 runway reconciliation: [paste each]. Draft a one-page roadmap
summary for an Investment Committee that states what the capital
buys, in what order, against which evidence, with which decision
gates, and within which runway. Map each capital tranche to the gates
it funds and the evidence that would release the next tranche. Flag
the single assumption the committee is most likely to challenge.
```

6 TROUBLESHOOTING

Symptom	Likely Cause	Fix
The roadmap is a list of ambitions with no evidence attached	It was built top-down before the I1 and I2 baseline was assembled	Return to Step 1. Build the evidence baseline from the channel scorecard, profitable-growth path, and runway, and rebuild every roadmap item on top of it.
Sales are scaling but new customers are churning	Growth outran delivery or onboarding capacity; the binding constraint was misdiagnosed	Re-run Step 2 against I2 cohort retention. If delivery is binding, add delivery capacity ahead of demand per Step 4.2 before pushing more volume.
Hires were made but growth did not improve	Headcount was added off the calendar, not against the binding constraint	Replace calendar hiring with the trigger-based capacity plan in Step 4. Tie every hire to a leading-indicator threshold and the constraint it relieves.
One channel got the whole budget and it saturated	The plan was a single big bet rather than a sequence of staged, reversible bets	Re-sequence per Step 5 into staged bets with gates. Cap any single bet and release the next only when its gate clears.
A bad quarter was only visible after it ended	The roadmap was tracked on lagging revenue with no leading indicators	Define leading indicators per Step 6.1 and wire them to the H2 dashboard so problems surface weeks ahead of the revenue miss.
The plan looks fundable but the venture nearly ran out of cash	The roadmap was never reconciled to burn and runway in the downside scenario	Run Step 6.2. Roll spend and revenue into C2 and confirm the downside scenario holds within runway; cut or defer the lowest-ranked bets until it does.
The Investment Committee cannot tell what the capital actually buys	The roadmap was presented as a narrative, not as tranches mapped to gates	Apply Step 7.2. Map each capital tranche to the specific gates it funds and the evidence that releases the next tranche.

VALIDATION STEPS

Confirm each of the following before declaring the scaling roadmap ready for the follow-on capital case:

- [ ] The evidence baseline is assembled from the I1 scorecard, the I2 profitable-growth path, and the C2 runway, with contradictions resolved
- [ ] The binding scaling constraint is named and tested against the evidence, not assumed
- [ ] Predictive baselines are built from historical data and are explainable enough to defend to the Investment Committee
- [ ] Candidate scale investments are ranked by expected return within payback and runway limits
- [ ] Base, upside, and downside scenarios are modeled and the downside still fits within runway
- [ ] Every capacity hire and investment is tied to a leading-indicator trigger and a named owner
- [ ] The roadmap is sequenced quarter by quarter with an evidence-based decision gate at each stage
- [ ] Leading indicators are defined, wired to the H2 dashboard, and reconciled to C2 burn and runway
- [ ] The scaling roadmap package is assembled and each capital tranche is mapped to the gates it funds; for CVC tracks, the PMS and J3 monitoring linkage is confirmed

8 NEXT STEPS

With the Scaling Roadmap complete, the Growth Acceleration and Unit Economics section of Phase Four is finished. The venture now has a repeatable motion (Guide I1), proven economics (Guide I2), and a sequenced, funded-ready plan to scale them (Guide I3).

Proceed to Section J, Follow-On Capital, Portfolio Value Creation and Corporate Integration. The Scaling Roadmap is the primary input to Guide J1, the follow-on funding and capital deployment case, where it is paired with the growth and economics evidence to build the investment memo.

Keep the roadmap live. As gates clear and constraints shift, update the sequence and the leading indicators on the H2 dashboard so the plan the board sees is always the current one, and so the leading indicators are ready to feed the automated portfolio-health monitoring in Guide J3.

MASTER CHECKLIST

A. EVIDENCE BASELINE

- ☐ Channel scorecard pulled from Guide I1 with verdicts and scalable ceilings
- ☐ Profitable-growth path pulled from Guide I2 with scale/fix/cut by segment and channel
- ☐ Burn and runway pulled from Guide C2 and current
- ☐ Contradictions between the three sources resolved

B. BINDING CONSTRAINT

- ☐ Each candidate constraint tested against the baseline evidence
- ☐ The single binding constraint named
- ☐ The next constraint anticipated for the following stage

C. PREDICTIVE PRIORITIZATION

- ☐ Explainable predictive baselines built from historical funnel, retention, and cost data
- ☐ Candidate investments ranked by expected return within payback and runway limits
- ☐ Base, upside, and downside scenarios modeled
- ☐ Downside scenario confirmed to fit within runway

D. CAPACITY PLAN

- ☐ Every capacity item tied to a leading-indicator trigger and threshold
- ☐ A named owner assigned to act on each trigger
- ☐ Delivery and support capacity sequenced ahead of the demand it serves
- ☐ Predictive indicators reused to feed the J3 monitoring

E. SEQUENCED ROADMAP

- ☐ Growth bets and capacity laid out quarter by quarter
- ☐ An evidence-based decision gate defined at every stage
- ☐ Highest-ranked, constraint-relieving moves front-loaded
- ☐ No single bet exceeds its cap; each stage is reversible

F. INDICATORS AND RUNWAY

- ☐ Leading indicators defined per bet and per constraint
- ☐ Roadmap reconciled to C2 burn and runway in the downside scenario
- ☐ Indicators and gate status wired to the H2 dashboard

G. ROADMAP PACKAGE

☐ Scaling roadmap package assembled for Guide J1

☐ Each capital tranche mapped to the gates it funds

☐ For CVC tracks: PMS and J3 monitoring linkage confirmed with the Close Owner

H. FINAL CHECK

☐ Every roadmap item is backed by evidence, triggered by an indicator, and reconciled to runway

☐ The capital request maps to staged gates, not a single bet

☐ The plan would survive Investment Committee scrutiny at Guide J1

If any bet is unbacked by evidence or breaches runway in the downside → It does not belong on the roadmap yet.