RIGHTTACK CONSULTING · PRODUCT STRATEGY & PORTFOLIO

The roadmap is what engineering was going to build anyway

Three sessions, three weeks apart, that end in decisions rather than a document — sizing the market by application, scoring every competitor on one scorecard, and finding the unit the customer actually buys in. What you keep is the structure, not the deck.

Ask a management team why the roadmap is the roadmap and you get a good answer for the first two items. After that it thins out fast. One item is there because a large customer asked loudly. One is there because an engineer has wanted to build it for three years and finally got it past a gate. Two are there because a competitor announced something and nobody wanted to be the person who said we should ignore it. And several are there because they were on last year's roadmap, and removing something from a roadmap requires an argument that nobody has time to have.

None of that is incompetence. It is what happens when there is no structure that a roadmap item has to survive in order to be on it. The plan is not wrong so much as untraceable — and untraceable means undefendable, which becomes a problem the first time a board, an acquirer or a new investor asks the question you have never had to answer in writing.

Three sessions, three weeks apart

The method is deliberately slow in one specific way. Three stages — environment, situation, strategy — each ending in a decision, each separated from the next by about three weeks. The sessions are for deciding. The gaps are where the work happens: someone goes and gets the number, calls the customers, builds the comparison. A strategy offsite that runs environment, situation and strategy in two consecutive days produces a document and no evidence, because there was never a window in which anyone could go and find any.

And no stage starts before the one above it has an answer. That sounds obvious and it is routinely violated — most commonly by starting on the roadmap, because the roadmap is the fun part and everyone already has opinions about it.

Environment: split by application, not by geography

Size the market, map the value chain end to end, and split the market by **application** rather than by geography or by customer size. Geography is a sales convenience — it describes how you are organised, not how anyone buys. Applications are where the buying decision is actually made, because the requirement, the competitor set, the price tolerance and the reason for switching are all properties of the application and only incidentally of where the plant happens to sit. A segmentation that mirrors your own org chart will tell you what you already believe.

Then deal with disruption honestly, which mostly means dealing with it consistently. Every industrial market has three or four candidate disruptions in play at any time — a materials change, a regulatory shift, a new entrant with a different cost base, a technology that is not quite ready. The failure mode is not ignoring them; it is arguing each one on its own terms, so that the one with the most articulate advocate in the room wins. Test all of them the same four ways, and they become comparable instead of debatable. Comparable is what lets you rank them and put a number against them.

Situation: one scorecard, or it is not a comparison

Score your own business against the market requirements you just defined, then score every competitor against that same scorecard. One scorecard. Not a scorecard per competitor built from whatever information happened to be available about each.

> **Give every competitor its own scorecard and you have not compared them. You have described them.**

This is the single most common defect in competitive work inside industrial businesses, and it is invisible while you are doing it. Each competitor profile looks careful. Each is built from a different mix of sources — a trade show conversation for one, a lost-deal debrief for another, a published spec sheet for a third — and each therefore emphasises whatever that source happened to cover. Put them side by side and the comparison is not between the competitors. It is between the sources. The discipline of forcing every one onto the same axes, and marking the cells where you genuinely do not know, is what turns a folder of profiles into a decision you can defend.

The value metric: the unit the customer actually buys in

Then the piece that changes the pricing conversation, and usually the roadmap with it: define the value metric. Not your feature list, and not your cost plus a margin — the unit in which the customer experiences the value, and in which their own finance team would recognise it.

In capital equipment this is rarely the purchase price. It is cost per unit of output, or yield, or uptime, or throughput per square metre of floor, or the labour hours the machine removes. The value-based approach here follows Michael Chase's work on capital equipment, and the practical consequence is significant: once the metric is defined, a roadmap item can be priced in the customer's terms rather than argued in yours. A feature that improves yield by a point is worth a computable amount to a named application. A feature that improves nothing measurable is worth an engineering argument, which is a different thing.

Define it, then test it before trusting it. Take the metric to three or four customers in the target application and check that it is the number they actually manage to. Value metrics fail in a specific and predictable way — the manufacturer picks the thing the machine is best at rather than the thing the customer is measured on, and the whole pricing case then rests on a unit the buyer does not use.

Strategy: objectives first, roadmap second

Only now does the roadmap get built, and it gets built in a fixed order. Financial objectives first — what this business has to deliver, and by when. Then the product and technology roadmap that delivers those objectives, with each item priced in value-metric terms against the competition. Then the risks that could break it, named and owned. Then twelve months of actions, each with a start, an end and a name against it.

> **A roadmap built before the objectives is a wish list with dates on it.**

The order is the whole discipline. Build the roadmap first and the objectives get reverse-engineered from it, which produces a document in which everything is consistent and nothing is tested. Build the objectives first and the roadmap has a bar to clear: every item has to answer what it moves, by how much, in whose money.

Who owns it, and what you actually keep

One named owner per workstream, throughout. The facilitator runs the sessions and owns none of the content — and that is not modesty, it is a structural requirement. A strategy that lives in the consultant's model leaves when the consultant leaves. A strategy that lives in a scorecard your own people built, populated with evidence your own people gathered, is still there to argue with next quarter.

> WHAT IS STILL ON THE TABLE TWELVE MONTHS LATER
>
> - **One page of ranked opportunities and threats**, with the evidence behind each rather than the advocacy.
> - **A value metric in the customer's own units**, tested with real customers before it was trusted.
> - **One scorecard, applied to every competitor** — including the cells marked unknown, which are a finding rather than an embarrassment.
> - **A roadmap priced in value-metric terms**, each item traceable to an objective and to a movement in the market.
> - **A named risk register** for what could break the roadmap, with owners.
> - **Twelve months of actions**, each with a start, an end and a name.

What this does downstream

Product strategy is the first decision in the operating sequence, and getting it wrong is expensive in a way that is hard to see at the time, because the cost shows up in the next four decisions rather than in this one. A portfolio nobody can defend produces a product definition nobody can freeze, because the requirements keep moving. That produces a footprint sized for a mix that was never settled, a supply base qualified for parts that get superseded, and a ramp for volumes nobody believes. Each of those looks like a failure of its own discipline. Several of them are this one, arriving late.

Where this is the wrong tool

If you have one product and one customer, you do not have a portfolio problem. You have a concentration problem, and it is a different conversation with a different urgency — no amount of segmentation work makes a single customer less than a single customer.

And if the roadmap is sound but nothing ships, the constraint is engineering capacity or change control, not strategy. Re-deciding what to build will not make the release land, and doing strategy work in that situation is a comfortable way to avoid a harder conversation about why the last three dates slipped.

That is the work I do.

WHERE THIS GOES NEXT

What to sell, and why →

The product strategy and portfolio method — three sessions, one scorecard, a tested value metric and twelve months of owned actions. Or start a conversation: rtackconsulting.com/contact