

Ways of Working: Executive Summary

The operating model on one page: who owns what, how much process each piece of work earns, and what gets measured.

Ask a product and engineering team who owns production at 2am and you will usually hear three names, or none. This operating model makes that question boring, and here it is on one page: eight principles you can test against your own organisation this week.

EVERY RESPONSIBILITY HAS ONE NAME AGAINST IT

Product owns the problem and the order of the queue. Engineering owns the solution and the system. Two different people hold them on purpose, because their daily argument is where someone challenges the scope and someone else prices the ambition. Quality, production out of hours and each person's growth have named owners too, and anyone can tell you who. "The team" owns nothing, and neither does "whoever picks it up".

WRITE THE DECISION DOWN WHEN YOU TAKE IT

The problem comes before the solution, the plan before the code, a decision record before anything someone will live with later, and a runbook before anything that can break at 2am. The tracker shows what is in flight, blocked and shipped, so status costs nobody a meeting. If a new joiner has to ask how to deploy, the documentation has failed them.

PROCESS HAS TO EARN ITS PLACE

Work gets a tier before it starts: feature, enhancement or initiative, each with its own discovery, artefacts and decision authority, so a copy change never meets the machinery of a new product line. Review is compulsory where the risk is real: authentication, payments, personal data, migrations, shared interfaces and anything hard to undo. Elsewhere the team skips it by prior written agreement.

DEFAULTS BEND; SILENT DRIFT BREAKS

Each process here exists because going without it costs more than following it, and by that test daily standups, story points and review before every merge are absent. A named person can override any default in writing, and the team reviews the override afterwards. Exceptions are healthy. Silent drift is the failure that matters.

SHIP SMALL, AND SHIP OFTEN

Engineers commit small changes straight to the main codebase, switched off behind feature flags until ready, with the automated test pipeline as the gate and production under thirty minutes away. Large work ships in increments, each with a place to stop. Teams of five and up plan in two-week cycles; teams of one to four pull work continuously, with a hard limit on work in progress. The discipline is identical, and only the ceremony changes.

MEASURE THE SYSTEM, NEVER THE INDIVIDUAL

Four numbers describe delivery: how often the team deploys, how long a change takes to reach production, how often a change causes a failure, and how fast the team recovers. Next to them sit planned against unplanned work, use of the ring-fenced engineering allocation, and each initiative's outcome against the case its sponsor made. Individual velocity, hours and story points stay off the dashboard: anyone can game them, and they tell you nothing about outcomes.

AI RAISES THE FLOOR; PEOPLE ADD JUDGEMENT

AI drafts problem statements, acceptance criteria, tests, documentation and code, and reviews every change in the pipeline. It decides and accepts nothing: a person owns the outcome of everything it touches, and each use case has a named owner, defined inputs and defined outputs. Treat it as a tireless junior with an excellent memory. It is still a junior.

SHRINK THE TEAM, KEEP THE RESPONSIBILITIES

The model targets 25 to 40 people and compresses without losing anything that matters. At ten, every responsibility still exists, held by fewer and more senior people. Responsibilities bundle into one person; they never disappear, and nobody loses track of who holds them. Roles split back out when the work triggers it, never because budget turned up.