ROLLOUT CHECKLIST

Shadow mode, before you cut over

Run the engine beside your dispatchers for four weeks and change nothing. You are not testing whether it works. You are finding the rules you never wrote down.

Before you start

- [ ] Write down who is not allowed to get what.

  Certifications, licences, regions, customers with a named account manager. This list is the whole project; everything else is preference.

- [ ] Pick one queue, not the whole operation.

  One trade, one region, one product area. A shadow run across everything produces an argument, not a result.

- [ ] Agree what a disagreement means, in advance.

  Decide now whether the engine differing from the dispatcher is a defect or a finding. Afterwards, it is a fight.

Week 1 - mirror

- [ ] Send every real job to the engine and throw the answer away.

  Dispatchers work exactly as they did. Nobody in the field sees anything change.

- [ ] Record both decisions and the reason for the engine's.

  The pair - theirs, ours, why - is the only artefact this phase produces.

- [ ] Watch for jobs the engine could not place at all.

  An empty candidate set is a missing rule or a missing person, and it is the cheapest bug you will ever find.

Weeks 2-3 - compare

- [ ] Sit with a dispatcher and read twenty disagreements together.

  Half will be a rule nobody wrote down. Write it down. The other half are worth arguing about.

- [ ] Separate "wrong" from "different".

  Two people could both have done the job. That is not a miss; it is what fairness looks like from close up.

- [ ] Check the spread, not just the fit.

  If one person wins most contests, the load penalty is wrong - not the matching.

Week 4 - cut over one slice

- [ ] Let the engine assign one queue for real, with a person able to override.

  Overrides are data. A week with none means nobody is looking.

- [ ] Keep the shadow running on everything else.

  You are not finished; you have one queue of evidence.

Stop and roll back if

the engine ever offers work to somebody who is not allowed to do it. Not "flags" - offers. That is a rule that did not make it into the configuration, and it outranks every other finding in the run.

Four weeks is the shape, not the law

A single queue with a settled rule set is often two weeks. An operation whose rules have never been written down anywhere has taken longer than four, and that is itself the finding - the time went into the list, not into the software.

Nothing in this checklist is specific to us. Run it against whatever you are evaluating, including something you wrote yourself.

A reference implementation to shadow against - 5xer.com/solutions/

Field-service dispatch in TypeScript - hard vetoes, capacity, fairness and a decision trace per job, MIT.

Yours whether or not you ever talk to us. Free to copy, print and circulate.