

A claimed AI saving clears three questions before it is a number.

Technology, operations and the numbers, read as one thing. Applied with rigour, only a small fraction of an AI wish list survives as a business case.

1

Can the business say what right looks like, in advance?

Precisely enough that something other than a person can tell when it is wrong.

Where it can, the check is already running and a person handles the exceptions. Where it cannot, headcount is a permanent operating cost, across functions that no roadmap removes.

2

Does the saving survive crossing a border, and is the model built for that market?

Regulation, labour law, service norms, language, cultural preference and cost base all move it. Every sign points to data sovereignty going from a diligence question to a hard constraint on what can be deployed.

Definability is not just a property of the task. It is a property of the market it sits in. The models with the most marketing behind them encode one market's defaults and may pass the exam in your language while being wrong about how the work is actually done in your specific geography.

3

Does it reach the P&L?

Consumption pricing on use nobody has measured. Checking sits in the function owning the work. Maintenance & security in technology, review and sign-off in legal.

All of this is incremental to an AI invoice, which by itself may be calculated differently at the next renewal.

Same business, different markets.

In mature markets, dealer ordering runs through sophisticated portals, with fewer logistical and pricing variations, so predictive automation has a far better chance. In rapidly growing ones the same order arrives by phone, WhatsApp, portal and in person.

What right looks like changes with the channel mix, and the automatable share is a fraction of the whole.

115 candidates. Three went to production.

Network energy. Sites drop below 20% utilisation on a predictable pattern, so right is a written threshold. Network equipment is around 70% of energy draw and energy is 20 to 30% of OPEX, so the saving lands in the largest controllable line.

Fraud and identity. Detection and alerting already worked and the pane was already built. What cleared was summarising the case history and matching accounts where the identifiers do not reconcile, with the agent reading it before acting, so the check was already in the job and handle time gave the before and after.

Tier-1 support. Right is already scripted, and the cases that need judgement stay where they are.

All three cleared the first question. One region, so the second did not move them. And each saving landed in a line someone already owned.

The same job is not the same job everywhere.



Can what right looks like be written down in advance?

The bespoke build is mine. Tech DD to Value Creation delivery. Right was still decided case by case.

The same three questions apply to the software businesses & vendors selling this, and to the “AI-native” platforms being spun up now.

Blue: right can be written down in advance. Orange: it cannot. Placement is directional, not a benchmark.

First principles up front, or the hold period gets spent redoing work already paid for. Mechanics over models.

The same money and the same people either clean up, or they build something.