

Four questions before you put an agent in production

Companion page to my keynote at Portugal Digital Summit, 22 October 2026. If you were in the room, this is what stays after it empties. If you weren't, it is the talk in writing.

Hélder Dias Ribeiro · 22 October 2026

AGENTS, DATA AND COMMON SENSE

Four questions before you put an agent in production.

ABOUT THE DATA

1

Can it reach?

Its own identity, access to the systems it needs, and written rules on what it may change.

ABOUT THE DATA

2

Can it understand?

It reads what is there, and the meaning was agreed before it started reading.

YOUR DECISION

3

Who answers for it?

A named owner, and rules approved before it was built.

YOUR DECISION

4

Does it pay?

It recovers more than it costs to keep running, maintenance included.

HÉLDER DIAS RIBEIRO · PORTUGAL DIGITAL SUMMIT 2026

Over the last year we ran 55 AI solutions through our adoption programme. Of those, 47 made it to production. Eight did not. Only one stopped because of the model. The other seven stopped at one of four questions that we now ask every agent before we let it out of the pilot.

A pilot proves it works once. Production has to prove it every day. The four questions are the difference between the two.

The four questions

1 Can it reach?

Does the agent have an identity of its own and access to the systems it needs, with permission to read and to write, and written rules about what it may change?

Reaching went well with commercial contracts. The agent has its own identity, separates the signed contracts from the unsigned, prepares the follow-up and updates the control file itself: the one place it is allowed to write. It went badly with a non-core service we run with a partner: the data sat on their side, and an agent of ours does not get into systems we do not control. That part is still unresolved.

An agent needs an access badge too.

2 Can it understand?

Can it read what is there, and was the meaning agreed before it started reading?

This was the question that cost us the most: four of the eight stopped here. It went well with store inventory: an agent that combines fixed rules with AI reasoning to explain why a count does not match, and gets the explanation right first time in four out of five cases. It works because someone decided beforehand what counts as a discrepancy. It went badly with maintenance manuals: extensive, good documentation, all of it written for people. In a narrow scope it worked; when we widened it, it stopped. Our data platform had solved the tables. Documents written for people had never been data. Think of your own shared drive: how many years of decisions sit in presentations and minutes nobody will open again? Today, for every type of document an agent will read, we agree the structure first, and the new document is born in it. What was left behind is only converted when an agent is about to use it.

Nobody had a data problem until an agent needed to read it.

3 Who answers for it?

When the agent acts, is there a named owner who answers for it, with rules approved before it was built?

The same inventory agent has an owner, the Operations department, and the rule is written down: what it does not get right the first time goes to the store's stock lead. With the owner and the rules in place, the decision to scale it took one meeting, not ten. With a voice assistant meant to talk to customers, we got to the rules only after building it, and it never went live. What holds you back is not the rules. It is getting to them late.

An agent acts. A person answers for it.

4 Does it pay?

Does it recover more than it costs to keep running?

We measure in hours and in euros, and every solution has to pay for itself, maintenance included. An agent that wrote store visit reports worked, and recovered 20 hours a year. We did not adopt it: it would cost as much to maintain as one that recovers about 300 times more. Maintaining an agent means infrastructure, model usage and the people who review it when the documents, the systems or the model change. On our numbers, a year of maintenance runs between 30 and 40 per cent of the cost of building; in classic software, between 10 and 20. It depends on the model and on token prices, and it will move.

Building is paid once. Maintaining is paid forever.

A PILOT	PRODUCTION
Proves it works once.	Has to prove it every day.
• Hand-picked data gets it past the first two questions • Someone is always watching, so nobody has to answer for it • Nobody asks it to pay yet	1 Can it reach? 2 Can it understand? 3 Who answers for it? 4 Does it pay?
That is why so many pilots work and so few reach production.	Seven of our eight non-adopted solutions stopped at one of these.

Two are about data. Two are your decisions.

Reaching and understanding are solved with work: access, formats, agreed meaning. Answering and paying cannot be bought: someone has to decide who owns the agent, what the rules are, and whether it pays. A pilot gets around the first two with hand-picked data and skips the last two, because someone is always watching and nobody is asking it to pay yet. That is why so many pilots work and so few reach production.

The answers are built once

The questions are asked of every agent. The answers are built once, in an agent platform with one piece per question. Identity and access, so the next agent gets its badge the way the contracts agent did. Machine-readable formats and a semantic layer, so the meaning is agreed once instead of once per agent. A registry with an owner and rules enforced while the agent runs, which is what the inventory

agent had and the voice assistant did not. Cost and value per agent, maintenance included, which is what stopped the visit-report agent. Across all of it, every run is logged and evaluated: the result, the cost, and whether a new model does better than the last. Underneath, security and compliance, with a common part and a per-agent part; we designed that layer against the AI Act, ISO/IEC 42001 and the NIST AI RMF, as a reference, not a certificate. It is a first version, and it will change with the technology.

THE AGENT PLATFORM, FIRST VERSION

Questions per agent. Answers built once.

One piece per question. The owner and the go decision stay with the business.

1 · REACH

Identity and access

An identity per agent, permissions, an official entry point into each of our systems.

2 · UNDERSTAND

Formats and semantics

Machine-readable documents and a semantic layer: the meaning, agreed once.

3 · ANSWER

Agent registry

A named owner and rules enforced while the agent runs.

4 · PAY

Cost and value per agent

Hours and euros, maintenance included.

Every run is logged and evaluated the result, the cost, and whether a new model does better

Security and compliance: a common part, written once for all agents, and a part for each agent.

HÉLDER DIAS RIBEIRO · PORTUGAL DIGITAL SUMMIT 2026

What I would do on Monday

Take your most advanced pilot and write down the four answers, one line each. If any line stays blank, that is where it will stop once it leaves the pilot. The first two are work. The last two need whoever is in charge to sit down and decide, and that is the harder meeting.

We are only on version one. We will keep changing and improving our process and our platform, and we are available to discuss and learn with you. There is no need to learn this alone.

Hélder Dias Ribeiro, Chief Digital & Technology Officer, Sonae MC. ·
www.helderdiasribeiro.net/4Questions