PROSPYR 305 · FREE SELF-CHECK

TWELVE QUESTIONS

Is Your Agent Actually You?

A twelve-question self-check for any business where more than one agent can act.

A twelve-question self-check for any business where more than one agent can act. Prospyr 305 · free · no score, no email required to use it. Answer honestly; the uncomfortable ones are the point.

How to use this

Sit with whoever owns your systems and answer out loud. Every question has one of three answers: **yes**, **no**, or **I would have to check**. "I would have to check" is the finding — it is the same answer an auditor will hear.

Do not skip ahead. Question 7 is where most businesses discover something.

Part 1 — The cast list

1. Can you name every identity that can act on your systems? Humans, agents, scheduled jobs, service accounts, integration tokens, the API key in the old script. Not the categories — the list. If it lives in one person's head, write it down before you continue.

2. Does any agent sign in as a person? A shared login, a borrowed mailbox, an API token minted from a human's account. If the answer is yes, every action that agent takes will be recorded against that human, and the record is wrong by design.

3. How many files hold a live credential? Count, do not estimate. When we first counted ours, a single vendor key sat in twenty-one files. Yours is probably not one file either.

4. If one of those credentials leaked tomorrow, could you say which system used it? If two consumers hold the same key, the honest answer is no, and no log retention will change it.

Part 2 — The record

5. Take the last five consequential things your automation did. Money moved, mail sent, a record changed, a customer contacted. For each one, name the actor without asking a colleague. Five out of five, or note where you stalled.

6. Does your log record the model and agent version, or only the tool that ran? "An agent did it" is not an answer six months later, when the model has been replaced twice.

7. For every automated email or notification you send, can you prove it landed, and where? Not "the job exited zero." Not "the log says sent." Prove it arrived, at a destination that exists. We once watched a scheduled job ask a question for six weeks into a channel that had been deleted. The log was correct. The audit trail was correct. Nobody was ever asked.

8. When an agent's output was wrong, could you reconstruct the input it saw? Prompt, context, retrieved documents, tool results — the whole thing, not the summary it wrote.

9. Is there a written answer to "who approved this?" for every action that moves money? A name, a timestamp, and what they were shown when they approved.

Part 3 — The morning after

10. If a customer said "your system did this," how would you establish what actually happened? Answer in steps, not in adjectives. If step one is "ask the person who set it up," that is the defect.

11. Does every alert have an owner who is a person? Not a channel. Not a team. A named human who would notice its absence. A monitor nobody owns is decoration.

12. Would your answer to an auditor be a name, or a shrug? The renewal questionnaire wants an agent inventory with a named owner per agent. Answer this one as if the underwriter were in the room.

Reading your answers

- **Three or more "I would have to check"** — you do not have attribution. You have activity. That is normal and it is fixable; it is also the exact condition that makes an incident unreconstructable.
- **Question 2 is yes** — stop. Every action that agent takes is filed under the wrong name, including the ones you would need in a dispute.
- **Question 7 is unsure** — your delivery has never been verified, only attempted. This is the cheapest of the twelve to fix and the most embarrassing to explain.
- **All twelve clean** — genuinely, say so and keep the list current as you add actors. The list is the artifact, and it decays.

What we do about it

The **Agent Identity Audit** takes this from twelve questions to a dated scorecard: every actor on your systems, 10–20 real actions proved end to end, the shared credentials that make proof impossible, and the short list of changes that make the next action attributable.

From $5,000. If we cannot show you at least one consequential action you cannot currently attribute, you do not pay.

Prospyr 305 · West Palm Beach, FL · (561) 990-2428 · prospyr305.com