AI PLATFORMS

A second reading for every AI agent

The platform's log says what the agent did. The board asks whether that matches what was approved, and who can stop it. Those answers sit in two places that rarely meet.

Roltrader Consultancy Group · 7 min read



An AI agent platform usually comes with its own governance: policies, guardrails, an audit log. That layer answers a clear question: what did this agent do, inside the settings it was given. It does not answer the next questions, and those are the ones a customer's board, auditor and insurer ask.

Was the policy approved, and by whom? Does it match what was agreed with the business? Who can stop this agent, and has anyone checked that they can? Would the record hold up if someone else looked at it? Those answers sit partly with the customer and partly with the platform, and today they are assembled by hand, customer by customer.

Two sets of facts that rarely meet

The platform holds records: which agents exist, who changed their settings and when, which keys and service accounts were created, how often each agent ran. The customer holds a different set: a register of the agents it believes it has, the owner it named for each, the conditions it set before an agent may do more, and counts from its own systems.

Each set looks complete from where it sits. The gaps show only when they are read side by side: an agent in the platform's records that is not in the customer's register, a change with no owner recorded, two counts of the same thing that disagree.

> The first useful finding is a mismatch between two lists.

A reading

A second reading takes the records the platform already holds and sets them beside what the customer declares and counts. It produces counts and lists, in plain sentences, with the contradictions first. The owner of each agent decides what to do with it.

WHAT A READING CONTAINS

- What lines up, and what differs, between the platform's records and the customer's own facts.
- Where the records are silent, a line that reads Not captured.
- For each open line, who can close it: the customer, the platform, or both.
- For each agent, the conditions the customer set before it may do more, and which of them have evidence.
- A receipt whose digest can be checked offline, and a short form a board can read in minutes.

Six potential use cases

The easiest way to see what a reading does is to follow it through a situation. These six are potential situations of the kind a reading is used in, described without naming any company or customer.

An account that AI assistants pay from. Imagine a new account that lets people authorise AI assistants to pay on their behalf, within limits they set, with a sponsor bank behind it. The account holder wants to know what the assistant paid and whether it stayed inside the limit. The provider wants to see payments from an assistant that was not on a verified path. The bank wants an exceptions register: payments above a limit with no linked approval, and stated limits the records do not recompute. One reading serves all three, each with its own view.

A platform whose agents contact people. Imagine an AI platform whose agents call and message prospects for a company's sales team. The customer's risk team asks who set each agent up, who changed it and who can stop it. The reading sets the agents seen in the platform's records against the customer's register: declared but not seen, seen but not declared. It lists

changes with no recorded owner and shows whether stop checks and sample reviews have dated evidence.

From pilot to production. Imagine a bank that has run an AI agent as a pilot and must decide whether it may do more. The owner set the conditions beforehand: each action carries an approver, someone can stop the agent, a sample of its actions is reviewed on a date. For each condition the reading shows whether the evidence supports it, contradicts it, or is silent. Silent reads Not captured.

A platform beside the client's own register. Imagine a large insurer that keeps its own register of AI agents and counts from its own systems, and runs its agents on one enterprise platform. The reading sets the platform's agents against the register, shows counts of the same unit that differ between the two, and lists keys and service accounts created or deleted in the period. The risk lead takes the contradictions first to the agent owners.

An accountable role for each agent. Imagine a regulated lender that names a senior role as accountable for each AI agent and has to explain that to its board and supervisors. For each agent the reading shows the accountable role as declared, beside the evidence behind it. The board pack is evidence for the lender's own judgement.

One platform, many customers. Imagine an AI platform that is asked the same evidence questions by customer after customer. Each customer gets a reading of their own setup. The platform sees the lines it can help close, and adds a missing field once for everyone.

Lines close on evidence

A reading is most useful the second time. A later reading is compared with the earlier one, and a line is shown as closed because the later records show it. In

every one of these situations the same thing happens: the platform answers the lines it can close, the customer answers the rest, and the next reading shows what moved.

The reading works from identifiers, flags and counts.

Where RTSightglass fits. RTSightglass is a separate product designed to attach to an AI platform. It turns the platform's records, set beside the client's own declarations and counts, into one dated evidence pack of counts and lists.

About RTSightglass → · **The six use cases →** · **For AI platforms →**

Evidence you can check, from your own data.

© 2026 Roltrader Consultancy Group · Privacy