PROOF

Reporting is getting cheap. Proof is not.

When AI writes the status report, who checks it?

Roltrader Consultancy Group · 4 min read · first published on LinkedIn, 7 November 2026



AI already writes status reports. Since May 2025, Microsoft's Planner has had a Project Manager agent that builds a report from plan progress, milestones and risks in one click. In 2019, Gartner predicted that 80 percent of today's project management tasks would be eliminated by 2030 as AI takes over data collection,

tracking and reporting. It said tasks, not people, and seven years on that part reads as accurate.

What it leaves open is simple. When an agent writes the report from the same data the team entered, you get a restatement in better prose. It can be the most persuasive document on a programme and the least tested. A sponsor still asks three things that no fluent report answers on its own.

Is it true?

Orchestra computes the colour for schedule, budget, risk and team from the programme's own data and reads Unknown where there is none. Its Steering Signals set the team's own words, quoted verbatim, against the status reported. In PRAIS, Programme Recovery reads the programme's own definition of green, amber and red against the evidence, shows the colour that definition gives beside the one reported, and never defaults to green. Assure checks AI-delivered work against the client's own source with five tests, including whether every "complete" is backed by a real artifact, and ties an agent's reported output to the client's own totals.

Who decided, and on what basis?

PRAIS holds the independent baseline and records how each figure was measured, so the same measurement can be re-run after go-live. Orchestra's decision log holds each material decision with who made it and when to look at it again, and a correction supersedes, never overwrites. Sustain keeps the owners after go-live. The Independent AI Audit shows what AI exists across the business, who owns it and what data it touches, and finds the AI nobody registered.

Why did we slip?

Programme Recovery keeps a neutral, dated record of the baseline, every change, every dependency and every decision. It quantifies a driver only where the data carries an amount, shows the rest as unattributed, and blames no one. Assure then compares the benefit promised with the benefit that landed.

Why now

Autonomy is arriving by default. Gartner expects 33 percent of enterprise software applications to include agentic AI by 2028, up from under 1 percent in 2024, and at least 15 percent of day-to-day work decisions to be made autonomously through it. Oversight is behind: a Dataiku survey of 685 CIOs found 81 percent lack complete oversight of AI agents built outside IT, and an EY survey of 202 senior AI executives found 26 percent of those using agentic AI could not detect unauthorised agents operating inside the company. These are largely vendor-run surveys, so treat the numbers with care. The direction is consistent.

The reckoning has a date. Gartner expects over 40 percent of agentic AI projects to be cancelled by the end of 2027, citing escalating costs, unclear business value and inadequate risk controls. Those are the three things a programme must evidence.

Boards will be asked to declare. For companies under the UK Corporate Governance Code, Provision 29 requires a board declaration on the effectiveness of material controls, for financial years starting on or after 1 January 2026, with the first declarations in annual reports in early 2027. Where an agent acts on a material process, the board will want evidence it can point to. That is my reading, not the Code's text. In the EU, legislators agreed in 2026 to push the AI Act's high-risk obligations to December 2027 and August 2028.

The date matters less than the lead time: evidence takes quarters to build, not weeks.

Waiting is the expensive option. The baseline cannot be rebuilt afterwards. Whether a promised benefit landed can only be judged against a baseline set independently before the work started, and a decision trail, a record of what changed and why, and a view of which AI exists are cheapest before agents scale. That is my judgement, and it is the reason to start now.

The first month

The programme's definition of green read against its evidence. The AI in use listed with an owner for each. The baseline held. The material decisions on record. One agent that acts on a material process tested against the client's own totals.

Why this approach

It is independent of whoever delivers or sells the work. It is deterministic: the same inputs give the same answer, with no language model in the verdict. It is honest about gaps: what was not captured reads Not captured, never zero. It can run inside the client's own environment. And it is built for the advisers and delivery teams who run programmes, so their expertise leads and the suite supplies the evidence.

What it does not do

It reads what the client captures, at a point in time; it is not live monitoring. It does not read a model's reasoning. And it does not replace the judgement of the people accountable. It is what their judgement stands on.

Roles change. The need for proof only grows. If you sponsor or run programmes with AI in them, this is the conversation to have before the next steering committee.

To start it, contact Roltrader at roltrader.com, or one of the independent advisers who work with the suite. They disclose their relationship with us, and their expertise leads.

Sources: Gartner, 20 March 2019; Gartner, 25 June 2025, and its 2024 technology trends; Microsoft Tech Community on Planner status reports; Dataiku and EY surveys as reported; FRC, UK Corporate Governance Code 2024, Provision 29; EU Digital Omnibus on AI, 2026.

Evidence you can check, from your own data.

© 2026 Roltrader Consultancy Group · Privacy