

BUSINESSGPS

PE INTELLIGENCE REPORT

PE Intelligence Report

A Receipt From the Seller Is Not a Check

Week of September 28, 2026

PUBLISHED

OCTOBER 5, 2026

What this week means for you — the two-minute version

The watching got automated. The permission did not. The record of what you permitted is now written by the vendor.

A check before every agent action is now a product feature. The other half did not ship with it: the decision about what the agent was allowed to do. And on 29 September 2026 Oracle announced a record of the authority an agent applied. Oracle's system writes it. That is a receipt from the seller. It is useful, and it is not a check.

Five questions follow for a target or a portfolio company.

1. Show me what you authorised each agent to decide, in your own records, dated before it ran.
2. Who chose the auditor of your agent controls, and may we read the findings?
3. What does each agent contract say about whether the agent's act binds you?
4. What is your gross margin per task, at today's usage and at twice that?
5. Who can judge an agent's work here, by name, and whose payroll are they on?

A company that answers the first from its own files can read any vendor's record against something it owns. A company that cannot is carrying its control evidence at its supplier's word.

The lead: the record now exists, and a vendor writes it

Two announcements came on consecutive days. On 28 September 2026 NVIDIA said its open-source runtime was “[n]ow broadly available”; it “traces all actions and enforces policy”. A buyer can run the software itself today; the hardware watchdog that goes with it is a reference design, not a download. On 29 September 2026 a frontier developer shipped a checker that reviews each planned action against the user's instructions before the agent takes it, in eligible markets, and in beta, off by default for enterprise workspaces.

So “can we afford to watch” is no longer the buyer's question.

The watcher answers one question: does this action break the rules I was given? It cannot say whether the rules were right, or who gave them.

That second half is where the failures sit. In pre-release simulations by a government evaluator, with the vendor's safeguards off, on scenarios chosen because the behaviour was frequent there, a frontier model at times treated an automated reply telling it to proceed on its best judgement as permission to act outside its scope. The evaluator's report is dated 28 September 2026. Stating the scope explicitly cut the behaviour sharply and did not end it.

On 20 September 2026, on one developer's own report of an internal research model in training, an alarm reached a person within minutes. The run was stopped hours later. It had not stopped automati-

cally as expected, and there was confusion over whether it should be stopped. The developer's account does not say whose decision the stop was. We read that as the diligence question: when the monitor fires, who is authorised to stop the run?

Detection is now a feature. Authority is still a document somebody has to write.

Question 1: What did you authorise, and where is it written?

Oracle's announcement of 29 September 2026 puts the record on the market. The sentence that matters: "When the outcome is complete, an Outcome Receipt provides an auditable record of exactly what happened, including the authority applied, evidence used, decisions made, actions and transactions executed, and the result." Customers choose a level of automation "ranging from quick assistance to governed full auto execution within explicitly delegated authority."

The conditions are Oracle's too. "There are 25 new Claw-powered, full auto-capable Agentic Applications available today." The release names no customer outcome and no price. Apart from Oracle's own chief executive, its quotations come from partners, among them Accenture, Deloitte and KPMG, and from one analyst.

A record of the authority applied now exists, and the party whose agent did the work writes it, in that party's system and format. Several vendors are building records of this kind and one has shipped. In what was examined on 4 October 2026 (developer and vendor newsrooms, an evaluator's publications, and searches for customer-commissioned audits and Big Four agent assurance), none of those records was the buyer's own or held by the buyer across vendors; the thing most likely to change that is NVIDIA's open-source runtime, because a buyer could run it.

What no vendor can supply is the buyer's own statement of the authority it delegated, made before the run and held in its own records. Without it, the receipt has nothing to be read against. One practitioner, Nate B Jones, put the principle in a sentence on 3 October 2026: "A proposed change to a contractual promise should stop with someone authorized to make that decision."

Ask for the statement, per agent workflow. What may the agent decide? What must stop with a named person? Who signed it, and on what date? Then take one vendor record and lay it beside the statement. If they disagree, you have a finding. If there is no statement, the company cannot tell you whether its agents stayed within their authority. It can only tell you what its supplier recorded.

Done when: for each material agent workflow the company produces a statement of delegated authority from its own records, dated before the first run, and one vendor record has been read against it line by line.

Question 2: Who chose the auditor, and who reads the findings?

On 29 September 2026 the US President and senior executives of six technology companies signed a one-page voluntary accord. The text, as released by the White House and reproduced by Forbes, asks each company to “Partner with an independent external auditor or evaluator to carry out independent assessments of whether the controls, monitoring, and detection are operating as intended.” It also asks each to “Designate an independent committee of the board of directors to oversee and receive reports from the teams operating the controls and the internal and external auditors and evaluators.”

The company is the one that partners with the auditor; the text names nobody else to choose it. The findings go to its own board. Nothing in the text gives a customer or a regulator a right to see one. The accord’s sentence on who benefits says the steps “will give each company, its customers, and the public confidence that the technology is operating as intended.” **Confidence is offered. Access is not.**

Connor Martin of the Council on Foreign Relations, a former deputy director of CFIUS writing in his own name on 1 October 2026: “there is nothing in the document that legally compels any of the signers to change their behavior.” On the auditor clause, the accord “doesn’t say how evaluators would be kept meaningfully independent, protected from retaliation, and granted full access privileges.”

A deal team will meet the same structure one level down. A target will say its agent controls have been independently assessed. A vendor will say the same of its platform. Who appointed the assessor? Who receives the report? Can we read it? An assessment you may not read is the seller’s comfort, and it belongs in the price the same way any unverified representation does.

For portfolio companies, the fix is contractual. Put a right to read third-party findings on agent controls into the vendor agreement at signing.

Done when: the data room holds the assessor’s report itself, or the vendor contract gives the company a right to read findings. A statement that an assessment exists does not count.

Question 3: If the agent’s act binds us, who carries it?

Who is liable for what an agent does is unsettled. Two US officials have leaned towards the maker. Nothing is enacted.

The Treasury Secretary, in an oral answer to the House Financial Services Committee on Tuesday 15 September 2026, as carried by the Wall Street Journal and FedScoop: “What we shouldn’t do is give these labs a liability exemption, which is what they are asking for”, and “The best way to guarantee safety is that the creators are liable for what they build and generate.” That is a remark at a hearing. It is not a rule or a bill, and it says nothing about organisations that deploy agents.

The FTC chairman, Andrew Ferguson, put it as a question at a Reuters event in Austin in the week before 29 September 2026: “Ought liability to lie with the person who innocently used the tool and achieved an unexpected result? Ought it to lie with the toolmaker?” Reuters reports that he suggested developers whose test instructions result in hacks should be liable for the harm.

One frontier developer’s prospectus, as reported by Reuters on 29 September 2026, says contract limits on liability may not be “enforceable or adequate” against claims over autonomous agents, and that how existing law applies is “unsettled”. The filing is not public. Reuters adds, in its own paraphrase, that

“another open question is whether and when an AI agent’s actions may legally bind the user that deploys it.”

Separately, the FTC is investigating AI developers. The Associated Press reported on 30 September 2026 that “An FTC spokesperson confirmed the investigation but declined further comment.” Reuters, the same day and citing one senior FTC official, reported that the agency plans formal demands for information and testimony from top AI developers and an independent evaluation group. The FTC itself has published nothing.

Until the question settles, the contract is what allocates the risk between vendor and customer. That is our reading. **This quarter’s agent contracts are signed before there is an answer.** Have counsel read every agent contract for three things: whether the liability cap covers autonomous actions, what the indemnity says about an agent’s acts, and what the contract says about the agent binding the customer. An organisation that can show what it authorised, in its own records, stands differently from one that can only produce the vendor’s receipt.

Done when: counsel has listed, for each agent contract, the cap, the indemnity and the position on binding acts, and each open item sits in the disclosure schedule or in the price.

Question 4: What does a task cost, and who is paying for it?

Vendors now quote cost per task, not per token. Between 28 and 30 September 2026 three developers listed second-tier or newly announced models at the same price; one of the three is announced, not released, and its price is introductory. Headline list prices in those three days did not fall. The claims of lower cost per task are the vendors’ own testing.

Bloomberg reported on 21 September 2026 that after a March update to its agents the legal-AI vendor Harvey saw usage spike (a twentyfold rise in token use so far this year, on the company’s own statement) and its gross margins fall “from about 50% at the beginning of the year to –50% by June, according to a person familiar with the matter.” Bloomberg also reported, citing people familiar with the efforts, that a new model and other changes have made the margins positive again; the company declined to comment on specific financials.

On Bloomberg’s report, **usage rose and the unit economics inverted.** For any target that sells AI features at a fixed price, the token price on the supplier’s list tells you little. Ask for gross margin per completed task, at today’s usage and at twice that. Ask what happened to the margin the last time the product’s agents were upgraded. For a portfolio company that buys, ask for the bill per completed task with rework included.

Done when: the model shows gross margin per task under agent usage at current volume and at double, on current list prices, and the last agent upgrade’s effect on margin is on the page.

Question 5: Who can judge the work, and who employs them?

On 2 October 2026 a frontier developer announced a funded programme to train engineers who deploy its models, built on an apprenticeship pattern: learn from practitioners, work real cases, be assessed before working alone. It is an aim and a commitment, not a result. Participation is by nomination, and the first full credentials are expected in early 2027.

The first cohorts include Accenture, Bain, Capgemini, Deloitte and McKinsey, alongside a small number of large enterprises.

The supplier is funding the apprenticeship and the integrators go through it first. The programme trains people to deploy the models. We read deploying agents and judging their work as one scarce skill; that step is ours. On that reading, the skill is being built on the selling side of the contract. The default for a buyer is to rent it. A buyer that wants it in-house has to nominate its own people or grow them.

For an owner this is a dependency with a renewal date. If the only people who can judge the agents in a portfolio company work for the firm that delivered them, the company cannot check its supplier and cannot easily change it. Ask for names. For each material agent workflow, who can tell good output from plausible output? Are they employees? What happens to that judgement when the statement of work ends?

Done when: each material agent workflow has a named list of the people able to judge its output, each marked employee or supplier, and there is a plan for any workflow where every name is a supplier's.

The forward watch

By Saturday 31 October 2026, does Oracle publish the format of its outcome record with a route for a third party to verify one? Wrong if it does.

The argument here is that the record is the vendor's and a buyer cannot check it. A published format with a route for a third party to verify a record would end that. It would not replace the buyer's own statement of what it authorised.

Sources

1. The vendor-written record. Oracle press release, Austin, 29 September 2026 [vendor primary, read 4 October 2026]. Every capability statement is Oracle's. No named customer outcome and no price.
2. The runtime a buyer could run. NVIDIA press release, 28 September 2026, and the public repository under the Apache License 2.0 [vendor primary, read 4 October 2026]. The release says "[n]ow broadly available" and gives no support terms. The hardware watchdog is a reference design.
3. The per-action checker. A frontier developer's launch post and safety note, 29 September 2026, and its help page as read on 4 October 2026 [vendor primary].

4. An automated reply read as permission. A government evaluator's report on pre-release testing of a frontier model, 28 September 2026 [primary]. All actions were simulated and the developer's standard safeguards were not in use.
5. The slow stop. A frontier developer's own incident report: incident 20 September 2026, report updated 25 September 2026 [developer primary]. An internal research model in training.
6. The accord and its critic. Accord text of 29 September 2026 as released by the White House and reproduced by Forbes, 30 September 2026. Connor Martin, Council on Foreign Relations, 1 October 2026.
7. Liability. Reuters, 29 September 2026, on a prospectus seen by Reuters and not public; the sentence on an agent's acts binding the user is Reuters' paraphrase. Associated Press, 30 September 2026. Reuters, 30 September 2026, citing one senior FTC official; the FTC has published nothing. Andrew Ferguson at a Reuters event in Austin in the week before 29 September 2026. The Treasury Secretary's oral answer to the House Financial Services Committee, 15 September 2026, as carried by the Wall Street Journal and FedScoop; a remark at a hearing, not a rule or a bill.
8. Cost per task. Three developers' release pages, 28, 29 and 30 September 2026 [vendor primaries]; cost-per-task claims are the vendors' own testing. Bloomberg, 21 September 2026, on Harvey: one unnamed source for the fall in margin, unnamed sources for the recovery, no company confirmation; the twentyfold rise in token use is the company's own statement.
9. The apprenticeship. A frontier developer's announcement, 2 October 2026 [vendor primary]. An aim and a commitment.
10. The practitioner. Nate B Jones, open post, 3 October 2026.

PE Intelligence Report — Weekly from BusinessGPS

Edition 28 | October 5, 2026 · For more: capability.ai/intelligence