

POLYGLYPH ANALYTICA · LIBRARY · THOUGHT LEADERSHIP

Why AI Pilots Stall in a Record Funding Year: The Data Gap, Oct 2026

Leaders say they are more confident about AI, and funding is at record levels. So why do so many pilots stall? We read four reports side by side to find out.

28 OCTOBER 2026 · [POLYGLYPHANALYTICA.COM/LIBRARY](https://polyglyphanalytica.com/library)

The argument in brief

Four recent reports, read together, tell a story that none of them tells alone. CB Insights shows venture money arriving at record scale, concentrated in very large rounds. KPMG shows large US organisations growing more confident in how they govern AI. Iron Mountain and Confluent show that, underneath, most organisations cannot trace their AI's data or keep it fresh enough to trust.

Our argument is that these findings are compatible. Organisations have built oversight of what their AI does and spends faster than oversight of what it relies on. Confidence is rising because the first kind of oversight is easier to see and to buy. The second kind decides whether an AI system can be trusted to act.

That has consequences. For investors, it shifts where durable value is found. For operators, it changes the order in which to spend effort. For everyone, it means the headline confidence numbers are statements about how visible AI activity is. They say little about whether the information is ready.

A caution about method. These are surveys of different populations, run by organisations with their own commercial interests, and we have not seen the underlying data. We draw on what each source states, we say where we are inferring, and we mark our own view as ours.

What the capital flow implies

CB Insights published its State of Venture Q3'26 report on 6 October 2026. It reports \$698 billion of venture funding through the third quarter, which it says puts 2026 on pace for a record above \$930 billion and for a result that beats the 2021 annual record. Deal activity keeps shrinking while the money grows: average deal size more than doubled, and roughly 80% of third-quarter funding came from rounds of \$100 million or more, with Databricks, Helsing and Kalshi named among them. The US took 59% of funding on 42% of deals, Europe had a \$31.5 billion quarter, its strongest since 2022, and M&A hit a four-year low, down 22% quarter on quarter.

Three things follow from these figures. First, concentration changes behaviour. A company that has raised a very large round has a mandate to grow quickly, and growth in enterprise AI means customers deploying it. Second, a market with fewer, larger deals and weak M&A offers fewer ways out. A weak exit market raises the cost of a deployment that drags, because there is less chance of being bought before the numbers have to show. Third, the money is chasing capability. Very little of a mega-round buys the customer's data readiness, which is on the customer's side of the contract.

This is our inference. CB Insights does not make this claim. It matters because it says where the strain will show. If funded companies need customers to deploy quickly, and customers are slowed

by their own information, the friction shows up as long onboarding, stalled pilots and disappointing expansion. It would not show up as a failure of the models.

Venture funding, 2026 (US\$ billions)



Source: CB Insights, State of Venture Q3'26, 6 October 2026. The full-year figure is CB Insights' projection.

Where Q3 funding came from



■ Rounds of \$100M or more (roughly) ■ All other rounds (remainder)

Source: CB Insights, State of Venture Q3'26, 6 October 2026. CB Insights gives the first share as roughly 80%; the remainder is the difference.

US share of funding and of deals



Source: CB Insights, State of Venture Q3'26, 6 October 2026.

What confidence is measuring


KPMG's Quarterly AI Pulse, published on 24 September 2026, surveyed 314 US leaders at organisations with \$1 billion or more in revenue between 24 July and 25 August. Its picture is encouraging. Nearly six in ten leaders report measurable business value, most often from productivity (55%) and faster decisions (49%). Confidence in the governance and capabilities in place reached 73%, up from 57% the quarter before. Of the organisations surveyed, 62% are building, deploying or developing AI agents, up from 53%, and 25% are working on multi-agent systems, against 6% in the previous two quarters.

The governance measures are mostly about visibility. 74% now include cost reviews in AI approval, up from 61%. 70% use AI monitoring dashboards. 43% have usage or token budgets. 49% have defined high-risk use cases where agents may not decide on their own. These describe visibility of spend and activity, plus some limits on autonomy. They are good practice, and they are the layer a finance team and an IT team can see and buy.

One finding in the same report points the other way. The share of organisations building controls into their agents, alongside monitoring and evaluation procedures, fell to 30% from 43% two quarters earlier. KPMG says many are relying on guardrails already in place. Confidence is rising while purpose-built controls on agents are falling, which is what you would expect if confidence comes from the layer that already existed.

KPMG's report also covers asset management and private equity separately. There, 81% of leaders report measurable value, agent deployment rose to 30% from 19%, and three quarters do not allow agents to touch sensitive data without human oversight, up from 36% the quarter before. That last figure is a statement about authority, and it shows the instinct is right: where data is sensitive, people keep a hand on it.

What the governance measures cover (% of organisations)



Source: KPMG Quarterly AI Pulse Survey Q3 2026, 24 September 2026 (314 US leaders at organisations with \$1B+ revenue).

What the data layer says


Iron Mountain commissioned Vanson Bourne to survey organisations about AI and information, and SecurityBrief reported the findings on 20 October 2026. 86% of respondents had faced disruption to AI work in the past 12 months. 81% believe AI cannot scale beyond pilots unless siloed and unstructured content becomes AI-ready data, and only 39% are in what the report calls a digitized and structured group. Four in ten spend at least 41% of AI project timelines on manual data preparation. Disconnected systems and silos are the most common obstacle at 56%, ahead of legacy formats or physical archives at 53%. And 70% say they cannot effectively track the lifecycle of the data used in their AI models, including where it came from and how it moved.

Confluent, an IBM company, reported on 9 September 2026 from its 2026 data streaming report, a survey of about 4,600 IT leaders. 72% cited challenges around data lineage, timeliness and data quality assurance, and 69% said insufficient real-time data infrastructure was holding back their

ability to scale AI. Confluent's own framing is that the question has shifted from what models can do to whether the business can trust the data under them.

These two surveys, with different samples and different vendors, reach the same finding: lineage, timeliness and quality are the constraint. Iron Mountain measures whether data can be traced and governed. Confluent measures whether it is fresh and trustworthy. Neither is a measure of model quality.

What organisations reported about their data (% of respondents)



Source: Iron Mountain research by Vanson Bourne, as reported by SecurityBrief, 20 October 2026.

What IT leaders reported (% of respondents)



Source: Confluent 2026 data streaming report (about 4,600 IT leaders), as reported by SecurityBrief, 9 September 2026.

Reconciling the evidence

There are three ways to square rising confidence with a weak data trail.

The first is that the populations differ. KPMG surveyed large US organisations with \$1 billion or more in revenue. The coverage we read does not describe the Iron Mountain or Confluent samples as limited to large US organisations. Large organisations may well be better prepared. This is plausible and partly true, but it does not explain why confidence is rising fastest while purpose-built agent controls are falling in the very sample where confidence is highest.

The second is that confidence leads capability. Leaders feel better about governance because the visible parts, cost reviews and dashboards, have improved, and they expect the rest to follow. If so, the data findings are a lagging indicator, and confidence is a bet on catching up.

The third, which we favour, is that the surveys measure different layers. Spend, usage and monitoring are about what AI does. Lineage, freshness and ownership are about what AI relies on. Organisations have improved the first quickly because it is visible and fits existing budgets and committees. The second needs work across systems and departments that nobody owns, which is why silos top Iron Mountain's list of obstacles.

What would change our mind: evidence that the same organisations report both high confidence and strong lineage, or that purpose-built agent controls recover while the share who can trace their data stays low. No source here lets us test that, and we do not claim to have.

Two layers of oversight (our framing, no data)

	Oversight of what AI does	Oversight of what AI relies on
What it covers	Spend, usage, dashboards, limits on autonomy	Origin, ownership, freshness, permission
Who can see it	Finance and IT, through existing tools	Nobody in particular, across many systems
What the surveys show	Improving, confidence rising (KPMG)	Weak, most cannot trace it (Iron Mountain, Confluent)

From a data problem to an authority problem

While AI only answered questions, a gap in the data trail cost you accuracy. Once AI acts, the same gap costs you control. An agent that raises an order, changes a record or sends a message is exercising authority on the strength of what it believes. If nobody can show where that belief came from, nobody can say whether the authority was well used.

This is why we think about it as a ladder. The right amount of authority for an AI system depends on what you can show about its inputs. A system that reads and summarises needs little trail. A system that drafts needs the trail for review. A system that acts with approval needs the trail so the approver can see what it saw. A system that acts within limits on its own needs the trail, a way to detect stale or unauthorised inputs, and a log that stands up afterwards.

KPMG's finding that 49% of leaders have defined high-risk use cases where agents may not decide alone is the right instinct, and it is a rung on this ladder. Our point is that the rung is only as safe as

the data trail beneath it. A limit on what an agent may decide does not tell you whether it decided on current, authorised information.

The authority ladder (our framing, no data)

Reads and
summarises

>

Drafts for review

>

Acts with approval

>

Acts within limits

What investors should do

Price the onboarding drag. If the customer's data is the constraint, then time to deploy sets sales cycles and net revenue retention, more than model quality does. Ask for time from signature to production by customer, and how much of it was data preparation. Iron Mountain's finding that four in ten organisations spend at least 41% of AI project timelines on manual preparation suggests the answer is rarely small.

Look for companies that do the data work. In a market where silos and legacy formats block deployment, a vendor that includes connection, mapping and provenance in its offer has a structural advantage over one that assumes clean inputs. Be sceptical of demos that run on prepared data.

Test provenance directly. Ask a company to show a customer where one output came from, and what happens when an input is stale, missing or unauthorised. A product that can show the trail is easier for a customer to approve and to expand.

Read the exit environment. With M&A at a four-year low and funding concentrated in very large rounds, a company with a long deployment cycle has fewer ways to recover time. Weigh runway against the realistic time to production.

What operators should do

Separate the two kinds of oversight in your own organisation. List what you can see about AI activity today, such as spend, usage and dashboards. Then list what you can show about its inputs: origin, ownership, freshness and permission. The second list is usually shorter, and it is the one that limits how much authority you can safely give.

Start with one workflow and the records it depends on. Name an owner for each record, set how current it must be and say who may use it for AI. Connect only those records, since the evidence points to silos as the most common obstacle.

Make the trail visible to the person who approves. For every output, they should see which records were used and when. If that is not possible, keep the system at a rung where a person checks the work.

Move up the ladder only when the trail has held. Wait for a period in which an agent's inputs were traceable and its outputs were checked before giving it more authority.

Build governance into each project rather than running a separate governance programme. Confluent argues for this, and its argument is reasonable on its own terms, though it also sells the products that implement it.

Three scenarios, with signposts

These scenarios are our judgement. No source forecasts them.

In the first, the gap closes quickly. Funded vendors and large customers invest in lineage and connection together, the share who can trace their AI's data rises, and deployment cycles shorten. Signpost: the lifecycle-tracking figure that Iron Mountain reports moves well away from 70% in the next survey, and purpose-built agent controls recover in KPMG's series.

In the second, the gap persists and pilots keep stalling. Funding stays concentrated, deployment drags, and a wave of large rounds struggles to show returns. Signpost: deal counts keep falling while average deal size rises, M&A stays weak, and surveys keep reporting disruption to AI work at levels near the 86% reported here.

In the third, the gap is bridged by services and by a new layer of provenance tooling. Value shifts to companies that do the data work, and a class of vendors whose product is the trail grows. Signpost: customers ask for provenance in procurement, and funding rounds begin to name data readiness rather than capability as the use of proceeds.

We think the second is most likely over the next two quarters and the third over a longer period, with the first the least likely in the near term. That is a judgement, and the signposts above are how to check it.

A note on the sources

CB Insights sells data and research on private markets, and its figures are as stated in its report. KPMG is a professional services firm that advises on AI, and its survey covers large US organisations only. Iron Mountain sells information management services, and its framing, continuous lifecycle governance, matches what it sells. Confluent, an IBM company, sells data

streaming, and its recommendation to move governance to the source and use real-time data matches its products.

None of that makes the figures wrong. It means each source's recommendation should be read with its business in mind, and that the strongest evidence in this paper is where independent sources with different products point the same way. Here, two vendors who sell different things both find that lineage, timeliness and traceability limit AI scaling, and a consultancy's survey shows controls on agents falling while confidence rises.

Sources

1. CB Insights, State of Venture Q3'26, 6 October 2026
<https://www.cbinsights.com/research/report/state-of-venture-q326/>
2. KPMG, AI's Value Story Sharpens as Organizations Gain Confidence in Governance, Accountability and Workforce Adoption (Q3 2026 AI Pulse), 24 September 2026
<https://kpmg.com/us/en/media/news/q3-ai-pulse-2026.html>
3. SecurityBrief, Data issues stall AI projects at most organisations, 20 October 2026 (Iron Mountain and Vanson Bourne research)
<https://securitybrief.com.au/story/data-issues-stall-ai-projects-at-most-organisations>
4. SecurityBrief, Confluent: Stale, unreliable data could impair enterprise AI, 9 September 2026 (Confluent 2026 data streaming report)
<https://securitybrief.com.au/story/confluent-stale-unreliable-data-could-impair-enterprise-ai>

Disclaimers

This paper is general commentary published by Polyglyph Analytica for information only. It is not investment, legal, tax, accounting or other professional advice, and it is not an offer or invitation to buy or sell any security or to engage any service. Take advice from a suitably qualified adviser before you act on anything in it.

Every figure in the paper is taken from the third-party sources listed above and is reported as that source states it. Some sources are surveys of particular groups, and their results may not apply to other organisations. We have not audited or independently verified those figures. Charts redraw only figures the sources state. Where a chart shows a remainder, it is the difference from the stated figure, and the chart says so.

Where the paper gives our view, suggestions or questions, it says so. Those are opinions, not facts, and they come from no source above.

Forecasts and projections, such as CB Insights' full-year funding projection, are estimates by their authors and may prove wrong. Past results do not guide future results.

Third-party names, reports and trade marks belong to their owners. They are used to identify the sources and do not mean the owners endorse or are connected with Polyglyph Analytica or this paper.

We have taken care to be accurate as at the publication date. We give no warranty that the paper is complete or current, and we accept no liability for loss arising from reliance on it, to the extent the law allows.

© Polyglyph Analytica. All rights reserved. You may share the paper in full and unaltered. Please ask us before you quote it at length or reproduce its charts.