Open Future Forum

OPEN FUTURE FORUM

Investor AI Report

NOVEMBER 2026

November 2026 | Updated data note and operating report | Murray Newlands | Open Future Forum

Prepared 2 November 2026. Evidence cutoff: 31 October 2026, 23:59 America/Los_Angeles. This edition addresses how a portfolio company repairs, qualifies or withdraws an unsupported AI claim after investment. Investment-thesis selection belongs to the companion VC and CVC report.

About this publication: Open Future Forum publishes research for executive and investor audiences. Murray Newlands is its founder. This report separates OFF survey observations, external research, author analysis and illustrative scenarios. Vendor-authored guidance is identified as such and does not constitute an independent product assessment. The collection and supporting methods are available through the OFF research index: https://openfutureforum.com/research/

4,221 substantive body words | 5,282 total words
Tables, FAQs, methods and references are excluded from the body count.

In this report

01 The decision in brief
02 Executive summary
03 What November adds
04 Evidence and sample boundaries
05 1. Decompose the claim before requesting more data
06 2. Resolve definitions without rewriting the history
07 3. Reconstruct the minimum evidence chain
08 4. Assign repair work to the person who can change the result
09 5. Close with support, qualification or withdrawal
10 6. Compare evidence states across companies without inventing performance comparability
11 External evidence and its limits
12 An illustrative first-100-days repair sequence
13 The next management conversation
14 Investor questions
15 Citable facts
16 Methods and limitations
17 References
18 Related research and citation
19 Editable working tool

The decision in brief

Turn an unresolved AI claim into a specific piece of company-owned work: identify the disputed assertion, the evidence missing, the decision it constrains and the person authorized to close it. Closure should mean that the claim is supported within a stated boundary, revised or withdrawn. Completing a dashboard field is insufficient if the underlying uncertainty remains unchanged.

Executive summary

An investor can request an AI evidence register and still receive a report that nobody can reproduce. The missing work may be a baseline, a cost reconciliation, an acceptance definition or an unresolved operating dependency. Repeatedly asking management for a better dashboard can add reporting effort without fixing the underlying problem. The useful next step is to define the evidence repair as an operating task with an accountable owner and an explicit closure condition.

The updated OFF investor data illustrates why reported impact needs interpretation. Better products was selected by 190 of 364 respondents, cutting costs by 141 and helping customers by 133. Fifty-one selected nothing measurable yet. Eight selected both a positive category and nothing measurable, leaving 321 respondents in the directly calculated positive union. Those combinations may reflect different companies or different meanings of measurable impact; the instrument does not resolve the explanation. They should prompt clarification rather than silent recoding. [INVIMPACT; INV_POSITIVE]

The ownership question has a separate base of 373, with 185 selecting the CEO, 97 the CIO or CTO, 43 finance, 72 individual teams and 62 too early to say. These are portfolio-level respondent views. They do not identify the person who can resolve a specific evidence problem inside an investee. A company needs a named process owner and a reviewer whose acceptance is appropriate to the claim. [INVOWN]

November's contribution is an evidence-repair workflow. It shows how to split a broad claim into testable assertions, reconcile conflicting definitions, reconstruct the minimum supporting record and decide whether the resulting claim is fit for its intended use. A worked first-100-days sequence demonstrates the process without implying that every company should follow the same calendar.

The October report already introduced an evidence register and first-100-days review. This edition does not relabel those ideas as new. It develops the harder operational question: what happens when a requested record is incomplete, contradictory or incapable of supporting the claim management has made? The answer should be a proportionate repair, a narrower claim or a documented withdrawal.

Investor reporting can divert scarce staff from serving customers. The proposed process therefore begins with the decision affected and reuses company records wherever possible. The company owns the work; the investor receives information within its rights. An evidence task that cannot change a relevant decision should be reconsidered before it becomes another recurring reporting obligation.

What November adds

The investor data has genuinely expanded, with cumulative bases of 373 for buying ownership and 364 for impact. The later source contributes 151 answers to each question before overlap is resolved. The resulting change is a cumulative sample update, not a measured improvement in the same portfolio companies. The companion statistics registry preserves the earlier source, later source and deduplicated result separately. [INVOWN; INVIMPACT]

The main operating addition is a closure standard. A claim can be supported for a defined population, revised to match the evidence, retained as an explicit hypothesis or withdrawn from the use for which it was proposed. Those outcomes are more informative than a binary complete field in a reporting spreadsheet. They tell management and the investor what can responsibly be said and what decision can proceed.

A second addition is the treatment of conflicting records. Finance may define benefit as realized expenditure reduction, while an operating team uses hours saved. Product may count a task complete when the system returns an answer, while customer success counts it complete only after acceptance. The repair process should preserve the original meanings, then establish the definition required for the specific decision. It should not force different facts into a convenient total.

A third addition is correction propagation. If a material claim is revised, the company should identify where the earlier version was used, including management reports, board materials and relevant investor communications. The appropriate correction should follow the company's responsibilities and disclosure process. Fixing the underlying spreadsheet while leaving the old number in circulation does not complete the work.

The report's first-100-days illustration is an application of these additions, not a new promise of universal implementation speed. The sequence ends with accepted evidence states and remaining constraints. It does not reproduce the October portfolio-triage article's ranking, reserve allocation or cross-company risk tiers. Its focus is the company-level work needed to make one material claim usable.

Evidence and sample boundaries

The OFF investor measures are selective registration responses about portfolio experience. The respondent is the unit of observation. Companies may overlap across respondents' portfolios, and the survey does not link answers to audited financial results, deployment records or individual investments. A positive selection is therefore not a verified return or a count of successful portfolio companies.

The ownership and impact questions have different completion bases and permit multiple selections. Contradictory impact combinations are retained. They cannot be resolved by assuming that a positive selection overrides nothing measurable, or vice versa. The directly calculated positive union and the none-only count answer different questions and are reported separately in the registry.

Operator measures are used only as contextual prompts. For example, the cost question asks about real-time visibility into AI running costs, with a base of 77. It does not verify complete workflow economics.

The genuinely matched production-operator cost result was already reported in October and is not a new November finding. [COST; OPERATOR_COST]

No qualitative prevalence is inferred from uncoded free-text fields. The report does not present private event conversations as attributable research or publish identifiable respondent information. The evidence tickets, calendar and company scenario below are original illustrations. They demonstrate a repair method without claiming that the method has been tested across the OFF population.

Exhibit 1. Keep reported impact and evidence status separate

Investor response measure	Count and base	What remains to be established at a company
Better products	190/364; 52%	Accepted result, relevant customer population and comparison
Cutting costs	141/364; 39%	Expense category, complete cost boundary and financial conversion
Helping customers	133/364; 37%	Customer outcome and evidence of its occurrence
Nothing measurable yet	51/364; 14%	Whether measurement is missing, benefit is absent or scope differs
Positive and nothing measurable both selected	8/364; count disclosed	Clarification; no explanation is imposed by the analyst

Exhibit note: Updated cumulative INVIMPACT, multiple selections. The last row is a derived intersection, not an additional response option. Results describe respondent views, not verified company outcomes.

1. Decompose the claim before requesting more data

A statement such as AI improved productivity can contain several assertions. The system may reduce task time, preserve quality, increase accepted output and create a financial benefit. Evidence for the first assertion does not automatically support the others. The operating team should split the statement into the components required for the decision, then identify which component is actually uncertain.

The intended use of the claim matters. An internal learning review may accept a directional estimate that would be inadequate for a material board commitment or an external financial statement. The company should state where the claim will be used and who is responsible for that use. This allows the evidence requirement to be proportionate without allowing a limited estimate to acquire a stronger meaning as it moves between documents.

Start with the narrowest assertion that can be checked. If management claims lower service cost, identify the service, workload, period and cost category. If it claims improved customer outcomes, identify the outcome and acceptance event. A request for all AI data is unlikely to resolve either question efficiently. A request for the records supporting a defined assertion gives the company a tractable task.

The team should preserve the original wording and its source. That record helps explain why the claim was made and where a misunderstanding arose. It also prevents the claim from changing silently during review. A narrower statement may be the correct final result, but the revision should be explicit so that people who relied on the earlier version can understand the difference.

The evidence ticket should name the decision constrained by the uncertainty. Examples include expanding a workflow, changing a staffing plan or including a benefit in a management forecast. If the claim is interesting but does not affect a current decision, the company may choose a lighter review or defer it. This keeps the process connected to operating priorities rather than the investor's desire for a comprehensive set of metrics.

The first output is therefore a clear problem statement, not a larger data room. It identifies what management wants to assert, what part lacks support, why that matters and what evidence could resolve it. That statement gives finance, operations, product and the investor a common task without requiring them to adopt the same meaning for every use of productivity or value.

2. Resolve definitions without rewriting the history

Conflicting numbers can arise because teams measure different things correctly. Operations may count completed tasks, product may count successful tool calls and finance may count invoiced services. The first repair step is to identify the unit and boundary of each measure. Treating the disagreement as a data-quality error before understanding those definitions can erase useful information.

A definition agreement should specify the eligible population, event, time period and exclusions required for the current decision. It should also preserve the original measures for their appropriate uses. Technical activity can remain useful for engineering even when accepted outcomes are required for a business review. The goal is a documented relationship between measures, not a single vocabulary imposed on every function.

Historical comparability needs special care. If the definition of completion changes, the company should either recalculate the prior period on a comparable basis or label the break. It should not present a continuous trend that combines incompatible definitions. A shorter comparable series is more useful than a longer series whose apparent improvement depends on an unnoticed measurement change.

Disagreement about financial benefit should be handled similarly. Hours saved, capacity redeployed, avoided hiring and realized spending reductions are distinct categories. A team may have valid evidence of one and insufficient evidence of another. The repair process should identify the supported category and the additional action required for conversion. It should not discard a useful operational result merely because it is not yet cash, or count it as cash because that is the category an investor prefers.

The definition should be approved by the function responsible for the decision. Finance should accept a financial category; the operating owner should accept the workload and service definition; product or engineering should explain technical event semantics. This is more useful than collecting generic sign-offs. Each reviewer should know what assertion they are accepting and what remains outside their review.

Once agreed, the definition should be versioned. The record should show when it became effective, which reports use it and why it changed. That small discipline makes future repair easier because the next reviewer

can distinguish a genuine change in performance from a change in measurement. It also prevents a disputed number from reappearing later under a slightly different label.

3. Reconstruct the minimum evidence chain

The company should identify the shortest chain of records that can support the assertion. For a delivery-cost claim, that may connect eligible tasks to accepted outcomes and relevant cost records. For an authority claim, it may connect the approved action to the effective permission and a demonstrated intervention. The chain depends on the claim; a universal evidence bundle can create unnecessary work while missing the decisive record.

Each link should have a source and a known limitation. A task log may omit manual work performed outside the product. A finance allocation may be an estimate rather than a direct observation. A customer acceptance record may be incomplete for one channel. These limitations should be visible so that the reviewer can decide whether the claim must be narrowed or whether additional evidence is necessary.

Reconstruction should not silently fill missing history with current assumptions. If review time was not measured during the relevant period, management can estimate it and label the estimate, or collect a new prospective sample. It cannot truthfully describe a retrospective guess as an observed historical result. The choice between estimation and new measurement should depend on the consequence of the decision and the feasibility of obtaining better evidence.

Sampling can be appropriate when full reconstruction is costly. The company should define the selection method, population and limitations before examining the result. A convenience sample of the easiest successful cases should not be presented as representative of the entire workflow. The review should include the exceptions most likely to affect the claim, particularly where a small number of costly failures can dominate the economics.

Evidence collection should respect access and retention limits. An investor may need a validated aggregate rather than raw customer records. A company can provide a scoped review result, a reconciliation or an appropriately authorized inspection. The method should fit the investor's rights and the company's obligations. More data is not automatically better evidence if it creates unnecessary exposure or cannot be interpreted reliably.

The result should be reproducible by an appropriate reviewer. That does not require every investor to rerun the company's systems. It means the company can explain the calculation, locate the relevant records and identify the judgment applied. A claim that depends entirely on one person's undocumented memory remains fragile even if that person is highly credible.

4. Assign repair work to the person who can change the result

Evidence problems often remain open because the named owner cannot resolve them. An investor-relations contact may compile the report but lack authority to change product instrumentation or financial definitions. The repair ticket should be assigned to the person responsible for the underlying process, with the reporting coordinator supporting the handoff. Accountability should follow the work required.

The ticket needs a concrete deliverable. Add cost visibility is too broad to close. Reconcile the defined service's accepted outcomes to the specified monthly cost categories is more useful. The deliverable should state the scope, source, expected review and decision it enables. A due date is meaningful only when the required dependencies and staff capacity are understood.

Dependencies should be recorded explicitly. Product may need engineering to capture an event, finance may need a supplier invoice and legal may need to approve access to a customer record. If the task depends on another team, the owner should identify the handoff and the consequence of delay. Otherwise, a sequence of reasonable individual tasks can remain stalled without anyone recognizing the shared constraint.

The investor's role is to support resolution within the agreed relationship. That can include helping management identify expertise, challenging an assumption or clarifying the evidence needed for an investor decision. It should not create a parallel operating command structure. Management remains responsible for running the company, and the investor's access and influence depend on the actual governance arrangements.

Repair should be proportionate to the next decision. A restricted expansion may require a targeted evaluation rather than a complete redesign of the reporting system. A material financial claim may require a more rigorous reconciliation. The company should avoid building a permanent reporting apparatus around a one-time question unless the recurring decision justifies it.

The backlog should be reviewed for usefulness as well as completion. Some tasks become obsolete because the company changes strategy or withdraws the claim. Others reveal that the original question was poorly framed. Closing those tasks with a documented reason is better than keeping them open indefinitely or producing an unnecessary artifact to satisfy a process. The objective is better decisions, with evidence work serving that objective.

5. Close with support, qualification or withdrawal

Closure should reflect what the evidence now permits the company to say. A supported claim names the relevant population, period and limitations. A qualified claim narrows the wording or scope. A hypothesis remains explicitly unproven and should not be used as an established result. A withdrawn claim is removed from the use for which it lacked support. These states make the outcome of the repair visible.

The acceptance reviewer should test the assertion against the evidence, not merely confirm that the requested file exists. A completed spreadsheet can contain unresolved assumptions. A dashboard can show

a precise number without a valid denominator. The reviewer should be able to explain why the evidence is sufficient for the intended decision and which limitations remain material.

Adverse findings should lead to a decision. If a benefit is lower than expected, management may revise the forecast, narrow the deployment or improve the process. If the evidence is unavailable, it may retain a restricted experiment while withholding a broader commitment. The repair process should not end with a report that everyone acknowledges but nobody uses.

Withdrawal should be treated as a legitimate outcome. A company that discovers an unsupported claim and corrects it has improved the quality of its operating record. Incentives that punish every revision can encourage management to preserve weak numbers. Investors should distinguish an honest correction from repeated unsupported assertions or failures to address known problems.

Corrections need to reach the appropriate users of the earlier claim. The company should identify where the number or statement appeared and determine the required update through its normal governance and disclosure process. The action may be an amended management report, a board clarification or another appropriate communication. This report does not prescribe legal disclosure duties; those depend on the facts and applicable requirements.

The closed ticket should retain the original assertion, the evidence reviewed, the final wording and the decision taken. That record helps future reviewers understand why a number changed and prevents the same disagreement from recurring. It also gives the investor a more useful view of management quality than a dashboard in which every item is green and no claim ever changes.

6. Compare evidence states across companies without inventing performance comparability

An investor may need to understand which companies can substantiate material claims, even when their businesses cannot be compared on a common productivity metric. The comparable object can be the evidence state: supported within scope, qualified, unresolved or withdrawn. That is different from ranking the companies' AI performance or treating the categories as a universal maturity scale.

The investor should preserve company-specific definitions alongside any common reporting field. Accepted resolution in a support business is not equivalent to accepted analysis in a research business. The field can record that each claim has a defined outcome and appropriate support, while the underlying measures remain distinct. This allows useful oversight without pretending that heterogeneous outcomes can be added together.

Changes in evidence state should be explained. Moving from unresolved to supported may reflect better measurement rather than improved operations. Moving from supported to qualified may reflect a broader workload or a newly discovered limitation. The investor should ask what changed and which decision is affected before interpreting the movement as success or deterioration.

The reporting burden should remain limited to material claims and agreed information rights. A fund may have different access across companies and ownership structures. Missing information should not

automatically be treated as poor operating practice when the investor is not entitled to receive it. The record should distinguish unavailable to the investor from unavailable within the company.

This report stops short of prescribing portfolio rankings, reserve decisions or valuation changes. Those topics were addressed in the earlier portfolio-triage work and require a broader investment judgment. The present contribution is the company-level repair process that makes a later portfolio review more credible. A fund can only interpret evidence states well if the underlying company has a meaningful way to close them.

The practical result is a better conversation with management. Instead of asking why the company has not achieved an arbitrary AI score, the investor can discuss the material assertion, the work needed to support it and the decision that follows. That conversation is more likely to improve the operating record while respecting the differences between companies.

External evidence and its limits

FinOps Foundation guidance links technical and business units in cost analysis. It is relevant when a repair task must reconcile consumption with a meaningful operating outcome. The guidance does not verify a company's allocation or determine whether a claimed benefit is financially realized. The company still needs its own records and an accepted definition. [E-FINOPS]

Anthropic's agent-evaluation guidance distinguishes the outcome of a task from the agent's narrative about completion. For evidence repair, this supports checking the relevant external result when a workflow takes action. It is an engineering principle within vendor-authored guidance, not proof that a particular product or evaluation process is reliable. [E-ANTHROPIC-EVAL]

NIST's voluntary AI Risk Management Framework provides a structure for governance and risk work. Its relevance here is the relationship between context, responsibilities, measurement and action. A company should not treat the use of framework terminology as closure of a material evidence problem. The actual assertion and supporting record remain the subject of review. [E-NIST]

METR's February 2026 update discusses limitations in its newer developer-productivity research design. It offers a useful example of a research team making measurement uncertainty explicit rather than presenting an unreliable current estimate. The analogy to company reporting is methodological: when a measure no longer supports the intended claim, the claim or method should change. It does not establish an effect for OFF respondents. [E-METR]

The sources support a disciplined distinction between measurement and conclusion. Technical activity, accepted outcomes, financial benefit and management judgment are related but separate. A repair process should identify where the connection is missing. It should not use an external framework or research result to fill a gap in the company's own evidence without making that inference explicit.

Some useful decisions must be made before a complete record can be reconstructed. The company can state the remaining uncertainty, restrict the scope of the decision and collect prospective evidence. It may also decide that the cost of reconstruction exceeds its value and withdraw the claim. Those are legitimate outcomes when they are transparent. The weak outcome is to preserve the original certainty while acknowledging privately that the supporting record cannot be produced.

Consider a company that cannot reproduce the workload definition behind its original cost claim. FinOps terminology alone cannot repair the missing denominator, and an evaluation framework cannot reconstruct a financial period. Management may need to qualify the historical statement and start a prospective measurement period with a stable definition. If the claim was used in a board paper, customer presentation or financing discussion, the repair owner should identify those dependent documents and arrange an appropriate correction. That propagation task is a management recommendation in this report. Its completion requires evidence that the revised meaning reached the people relying on the earlier statement, not simply a new spreadsheet in an internal folder.

Evidence comparison

Source	Contribution to this decision	Limit on interpretation
FinOps [E-FINOPS]	Connects consumption to a business unit	The company must supply its own allocation records
Anthropic [E-ANTHROPIC-EVAL]	Separates reported completion from external result	Engineering guidance does not verify a financial assertion
NIST [E-NIST]	Assigns context and responsibility	Framework vocabulary is not closure evidence
METR update [E-METR]	Shows why an unsupported estimate may need revision	Methodological analogy, not an effect estimate for OFF

An illustrative first-100-days repair sequence

Consider a hypothetical portfolio company reporting that an AI-assisted service reduced delivery effort by 30%. The investor's initial review finds that the figure comes from a small set of successful tasks, excludes customer review and uses a different completion definition from the monthly operating report. The company does not need another general dashboard request. It needs to determine what the figure can support and whether the broader service is improving.

During the first two weeks, management records the original assertion and its intended use in the expansion plan. The operating leader defines the eligible service population, finance identifies the relevant cost categories and product explains the existing completion events. The immediate decision is whether to expand to a larger customer group. The company keeps the current restricted service operating while it resolves the evidence required for expansion.

Over the following weeks, the team reconstructs what it can from existing records and identifies what cannot be recovered reliably. It does not invent historical review time. Instead, it establishes a prospective sample that includes routine and exception cases, with an agreed acceptance definition and a record of review effort. The sample design and limitations are documented before the results are used to support the claim.

Exhibit 2. Illustrative evidence-repair backlog

Assertion or problem	Accountable owner	Closure condition
Reported effort reduction excludes review	Operating leader with finance reviewer	Complete effort boundary or explicitly narrowed claim
Completion events differ between systems	Product owner	Versioned mapping to the accepted business outcome
Historical evidence is incomplete	Process owner	Supported reconstruction or labelled prospective measurement
Expansion assumes the result generalizes	Business-unit leader	Evidence for the proposed population or restricted expansion
Earlier number appears in board materials	Responsible executive	Appropriate correction through the company's review process

Exhibit note: Hypothetical company and proposed tasks. The calendar is illustrative; no universal completion period or measured success rate is asserted.

By the middle of the illustrative period, the team can distinguish an observed improvement in the restricted workflow from an unsupported claim about the whole service. Management revises the wording accordingly. If the prospective evidence supports expansion, it presents the relevant population and remaining limitations. If it does not, management can narrow the deployment or continue learning without repeating the original 30% as an established company-wide result.

The acceptance review then asks whether the revised claim is sufficient for the intended decision. Finance does not need to certify the model's internal behavior; it needs to accept the financial interpretation within its scope. The operating owner does not need to determine disclosure obligations; it needs to explain the service result and limitations. The responsible executive brings those judgments together for the expansion decision.

Exhibit 3. Closure record

Field	Required entry
Original assertion	Exact claim and where it was used
Evidence boundary	Population, period, definition and source limitations
Final state	Supported, qualified, hypothesis or withdrawn
Accepted wording	Statement the evidence now permits
Decision consequence	Expand, restrict, defer, revise forecast or other specified action
Correction follow-through	Reports and recipients requiring an appropriate update

Exhibit note: Proposed operating record. Categories describe evidence status, not company quality or a performance ranking.

By the end of the illustrative 100 days, the useful deliverable is a closed decision record and a smaller set of specific unresolved tasks. Some claims may be supported, some narrowed and some withdrawn. The investor should understand which assumptions remain relevant to the investment and what management is doing about them. A perfect-looking dashboard is unnecessary if the company can explain its evidence and act on what it learns.

The next management conversation

Choose one material AI claim that management cannot currently reproduce. Ask what decision depends on it, which assertion is uncertain and what evidence would resolve that uncertainty. Assign the work to the person who can change the underlying record, then identify who can accept the result for the intended use. This is more productive than requesting a broad improvement in AI reporting.

Treat a narrower supported claim as progress. A company that distinguishes capacity from cash, a pilot result from a broad rollout or a system message from an accepted outcome has improved its ability to make decisions. The investor should encourage that precision while remaining alert to unresolved assumptions that materially affect the investment case.

November's recommendation is to make evidence repair an operating process with a real endpoint. The endpoint can be support, qualification or withdrawal, followed by the appropriate decision and correction. That gives the company a defensible record and gives the investor a clearer view of what has been learned after capital was committed.

Investor questions

Is this the same as the VC and CVC report?

No. That report tests the investment thesis. This report addresses company-owned evidence repair after investment.

Does a positive investor response verify a company result?

No. It is a respondent's broad portfolio impression. Company evidence remains necessary. [INVIMPACT]

Should contradictory responses be deleted?

Not without a justified rule. The eight positive-plus-none combinations are retained and disclosed; their meaning is not inferred. [INV_POSITIVE]

Is the first-100-days concept new in November?

No. October already included it. November adds explicit repair, acceptance, qualification, withdrawal and correction steps.

Does a completed evidence ticket prove the original claim?

No. Closure can support, narrow or withdraw the claim. The accepted wording and decision consequence must be recorded.

Must investors receive raw customer data?

No. Evidence access should fit information rights, confidentiality obligations and the decision. Aggregates or scoped review may be appropriate.

Citable facts

- Investor ownership base: 373; impact base: 364, updated cumulative. [INVOWN; INVIMPACT]

- Positive impact union: 321/364, 88%; eight respondents also selected nothing measurable. [INV_POSITIVE]

- None-only impact responses: 43/364, 12%; distinct from all 51 nothing-measurable selections. [INV_POSITIVE]

Methods and limitations

Calculation-v1.1 removes byte-identical exports and invited-only records and selects the latest eligible nonblank response per normalized respondent identifier per question. The published October-compatible baseline is reproduced. New event sources and overlap are disclosed in the statistics registry. Registration timestamps are not answer-edit dates.

The data is selective, self-reported and not linked to audited company results. Multiple selections and contradictory combinations are retained. No portfolio-company prevalence, causal effect, return or universal evidence score is claimed. The repair process and 100-day scenario are original recommendations and illustrations, not measured outcomes.

References

[OFF-DATA] Open Future Forum. Executive AI Statistics, November 2026. INVOWN, INVIMPACT, INV_POSITIVE, COST and OPERATOR_COST.

[E-FINOPS] FinOps Foundation. Unit Economics capability. https://www.finops.org/framework/capabilities/unit-economics/

[E-ANTHROPIC-EVAL] Anthropic Engineering. Demystifying evals for AI agents. 9 January 2026. https://www.anthropic.com/research/demystifying-evals-for-ai-agents

[E-NIST] NIST. AI Risk Management Framework 1.0, current official resource page. https://airc.nist.gov/airmf-resources/airmf/

[E-METR] METR. We are Changing our Developer Productivity Experiment Design. 24 February 2026. https://metr.org/blog/2026-02-24-uplift-update/

[OFF-PRIOR] Open Future Forum. Investor AI Report, October 2026. https://openfutureforum.com/research/investor-ai-report-october-2026

[OFF-TRIAGE] Open Future Forum. How Should Investors Run Quarterly AI Triage Across a Portfolio? October 2026. https://openfutureforum.com/blog/quarterly-ai-portfolio-triage-investors

Related research and citation

The November VC and CVC report tests investment theses, CFO qualifies financial benefits, AI Leaders examines changed operating scope, and Private Equity follows benefits through ownership transition.

Suggested citation: Newlands, Murray. Investor AI Report, November 2026, Updated Data Note and Operating Report. Open Future Forum. Evidence cutoff 31 October 2026.

Open Future Forum research: https://openfutureforum.com/research/

Editable working tool

Company evidence repair register: download the editable CSV worksheet. The worksheet implements the report's recommendation; it is not a validated score, mandatory standard or completed assessment.