Open Future Forum

OPEN FUTURE FORUM

CEO AI Leverage Report

NOVEMBER 2026

November 2026 | Edition 5 | Research report | Murray Newlands | Open Future Forum

Prepared 2 November 2026. Evidence cutoff: 31 October 2026, 23:59 America/Los_Angeles. No new comparable CEO-specific survey was collected. Buying evidence is carried forward; title-based comparisons are a disclosed reanalysis of historical answers.

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/

4,453 substantive body words | 5,925 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 this edition adds
04 How to read the CEO evidence
05 1. Allocate attention before adding initiatives
06 2. Choose the operating model for the work
07 3. Make delegation conditional on the original promise
08 4. Use expectation differences to inspect assumptions
09 5. Sequence the portfolio around shared constraints
10 6. Keep workforce commitments credible
11 External research and the executive decision
12 Working application: decide the sequence, then delegate
13 The CEO's next agenda
14 Reader questions
15 Citable facts
16 Methods and classification note
17 References
18 Related research and citation
19 Editable working tool

The decision in brief

The CEO should decide where AI deserves scarce management attention, which commitments can be delegated and what changes require an executive decision. Purchase approval is only one part of that job. November's contribution is a method for allocating attention and making delegation conditional on evidence, so useful work can expand without leaving authority permanently undefined.

Executive summary

The CEO remains the most frequently named participant in AI purchase sign-off in the carried-forward OFF instrument: 227 of 467 respondents selected the CEO, or 49%. CFO or finance appears in 144 answers, technology leadership in 79, an individual business unit in 45 and no single owner yet in 60. Multiple selections are permitted. The finding describes perceived purchase authority, not the actual authority to release, expand or stop a production workflow. [SIGN]

That distinction matters because the CEO's scarce resource is often attention rather than the price of a pilot. An initiative can require scarce data expertise, legal review, process redesign and business sponsorship while appearing inexpensive in the software budget. Several individually attractive projects may compete for the same people or depend on the same unresolved operating decision. A portfolio review should expose those dependencies before authorizing more work.

This edition proposes contingent delegation. The CEO assigns an operating owner and a bounded decision right, then specifies which changes return the decision to an executive. The boundary may concern spending, data, customer exposure, autonomy, operating dependence or the assumptions supporting a benefit. Routine improvements can remain delegated while consequential changes receive deliberate review. The point is to reduce repeated escalation without turning a one-time approval into indefinite authority.

The report also revisits CEO and finance expectations using an explicit keyword classification. Seventy of 102 historical respondents classified as CEO/founder selected measurable return below six months, compared with 30 of 71 classified as finance. These are reclassified cross-sections, not a new CEO panel or a change in sentiment. The classification differs from October, so its results cannot be presented as movement from the earlier role cuts. [PAY_CEO/founder; PAY_Finance]

A formal review process can recreate the bottleneck it is meant to solve. The CEO should therefore reserve personal attention for conflicts that cross delegated boundaries: incompatible objectives, scarce shared capacity, material exposure and commitments that are difficult to reverse. The operating register and decision-rights matrix in this report are designed to surface those issues. They should connect existing functional evidence, not require the CEO to inspect every model change or maintain another technical inventory.

What this edition adds

The November exports contain new security, marketing and investor answers, but they do not add a comparable CEO-only instrument. CEO dinner rosters are not converted into survey findings. The common buying questions remain historical observations with their original response coverage. This report uses them to examine executive decisions; it does not claim that CEO sentiment changed during October or that the November publication date makes those answers new. [SIGN; FUND; PAY]

The first addition is an allocation method for management attention. October explained why purchase approval must connect to production accountability. November asks which production decisions actually require the CEO and which should remain delegated. It distinguishes routine improvement from changes that alter the company's commitment, exposure or operating model. This provides a more specific answer than asking the CEO to own AI in general.

The second addition is a dependency view of the portfolio. Projects can compete for the same integration team, customer data, approval capacity or executive sponsor. A proposed investment may be attractive in isolation but poorly sequenced within that portfolio. The report supplies a review that asks which constraint should be resolved first and what other work should wait, rather than treating every funded pilot as an independent opportunity.

The third addition is a transparent reanalysis of title-based expectations. The November keyword classifier is reproducible but broader than the published October classification. It includes founder language and a wider finance vocabulary, with a documented precedence rule for mixed titles. That changes the group bases. The new result is labeled recalculated historical evidence, and the earlier role cuts remain historical rather than being silently overwritten or joined into a trend.

Finally, the report develops a worked executive decision involving two initiatives with a shared operating dependency. Its purpose is to show how a CEO can choose sequencing, delegation and review conditions without pretending that a simple ROI ranking resolves every strategic question. It is an illustration, not a private executive story. The basic question of how to use time saved by AI was already discussed in Murray's November 1 newsletter; this report extends the discussion to conflicting commitments and decision rights.

How to read the CEO evidence

The main purchase-authority finding comes from 467 eligible answers across the common instrument. It is not a survey of 467 CEOs. Respondents hold different roles and describe their own understanding of purchasing authority. The sample is selective and self-reported, drawn from event registrations and applications. Invited-only records are excluded, while eligible statuses do not establish attendance.

The role comparisons use the job-title field attached to the selected answer. The classifier first identifies CEO, chief-executive and founder terms; then finance, CFO, controller and treasurer terms; then technology and investor terms. A mixed founder-and-finance title therefore enters the first group under this rule. The grouping is an analytical convention, not an independently verified executive biography or a respondent-selected role.

The return-timing bases are 102 for CEO/founder and 71 for finance. Both clear the report's descriptive headline floor, but sample size does not remove selection, classification or company-mix effects. The groups are not a matched set of CEOs and CFOs from the same companies. Their answers cannot establish disagreement within a company, a causal effect of role or a psychological optimism score.

Operator evidence provides separate context for the decisions a CEO should ask about. It is not attributed to the title-classified CEO group unless a valid match supports that statement. Recommendations and illustrations are labeled accordingly. The objective is to improve the questions asked of the company's own evidence, while preserving the limits of the OFF observations.

1. Allocate attention before adding initiatives

A software budget can hide the real cost of an AI initiative. A pilot may require a senior operations leader to redesign a process, an experienced engineer to resolve a data dependency and a legal team to assess a new customer commitment. These people may already be the limiting resource for several other projects. The CEO should ask what management capacity an initiative consumes before treating its low subscription cost as evidence that it is cheap to pursue.

The carried-forward sign-off result helps explain why this question belongs on the executive agenda. The CEO is mentioned in 227 of 467 answers, the largest category. Yet the instrument says nothing about the amount of executive time involved or whether that involvement improves outcomes. The finding supports examining the approval path; it does not prove that centralizing more decisions in the CEO's office is effective. [SIGN]

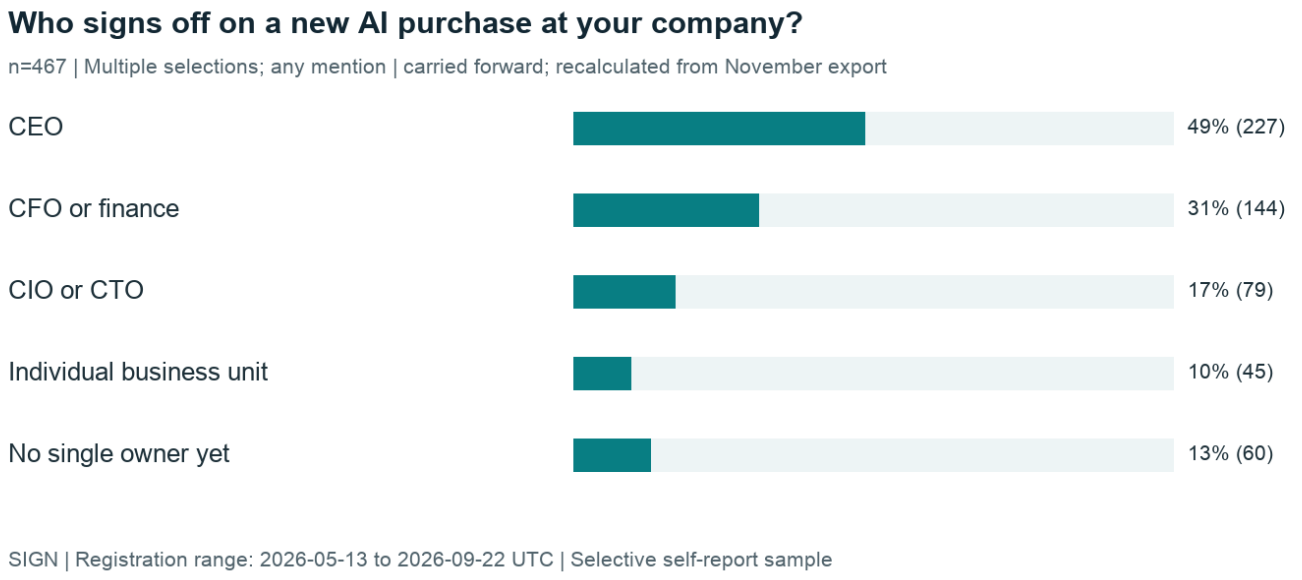
A useful attention allocation begins with the decision the organization cannot make locally. Perhaps two functions disagree about the value of a shared data project. Perhaps a service owner wants autonomy that security cannot support within the current design. Perhaps a promised efficiency benefit depends on a staffing or customer-service decision that no project team can authorize. These are executive issues because they cross authority boundaries, not because they involve AI.

Routine technical choices should stay with the teams equipped to make them. The CEO needs an escalation rule rather than a standing invitation to every design meeting. Ask the operating owner to identify the decision, options, consequence of delay and recommendation. A review that arrives with only a status update consumes attention without using executive authority to resolve anything.

The portfolio view should include the work that will be displaced. If an integration team can support only one material release this quarter, approving three projects does not create three delivery opportunities. It creates a queue, competing expectations and possible workarounds. The CEO should decide the order or add credible capacity, with the cost and delay of each option made explicit.

This changes the discussion from how many AI projects the company has to which constraints deserve executive intervention. An initiative that removes a shared bottleneck may deserve priority even if its standalone return is hard to express. Another initiative with an attractive forecast may need to wait because the organization cannot implement it responsibly yet. The decision should make that reasoning visible rather than allowing informal sponsorship strength to determine the sequence.

Exhibit 1. Reported purchase sign-off participants



Exact selection	Count	Base	Any mention
CEO	227	467	49%
CFO or finance	144	467	31%
CIO or CTO	79	467	17%
Individual business unit	45	467	10%
No single owner yet	60	467	13%

Exhibit note: SIGN; carried-forward common instrument, selective sample, multi-select. Purchase sign-off is not measured production-release, risk-acceptance or stop authority.

2. Choose the operating model for the work

The organization does not need one operating model for every AI workflow. A shared platform can provide common capabilities, while business teams remain accountable for the outcomes of their applications. Some activities may justify central operation because their data, control or service requirements are shared. Others may benefit from local experimentation within a narrow boundary. The CEO's decision should follow the work rather than an abstract preference for centralization or decentralization.

Start with what must remain consistent. Identity, data-use rules, spending controls and evidence requirements may need common standards. Customer experience, workflow design and the definition of an accepted result may vary by business unit. A centralized team can create reusable controls without taking responsibility for every commercial decision. Conversely, a business unit can own its outcome without inventing a separate security or procurement regime.

The failure mode is an unassigned handoff. A platform team may assume the business validates output quality; the business may assume the platform guarantees it. Finance may assume the sponsor captures savings; the sponsor may assume finance will find them in the budget. The CEO should require a clear division of responsibilities at those boundaries, including the evidence each team supplies to the next.

Operating-model choices also determine the cost of change. If every new use case requires a central committee to interpret the same basic rules, the model may not scale. If every team can expand authority without review, the company may lose the ability to understand its commitments. A reusable release process with explicit local discretion can reduce both problems, provided material changes have a clear route back to the appropriate owner.

The company should test the model on a real workflow before declaring a reorganization complete. Follow one request from business need through approval, deployment, monitoring and change. Record where the process stalls, where evidence is repeated and where no one can make a decision. Those observations reveal whether the operating model works better than a new organizational chart alone.

The CEO should then decide which handoffs warrant structural change and which need a clearer agreement. Not every coordination problem requires a new executive title or a new committee. Sometimes the missing element is a named business owner, a standard evidence record or a delegated budget limit. The operating model should solve a demonstrated decision problem, with enough flexibility to adapt as the company's actual use of AI changes.

3. Make delegation conditional on the original promise

A purchase approval often contains an implicit promise: the workflow will serve a particular purpose, use a particular set of resources and remain within an understood level of exposure. Over time, the application may gain users, connect to new data or take actions that were absent from the original proposal. The CEO needs a way to preserve useful delegation while recognizing when the promise has materially changed.

Contingent delegation starts with a clear scope. The operating owner can improve the workflow within agreed boundaries, provided the evidence remains adequate. The boundary should identify the kinds of change that require renewed approval, such as a new external commitment, materially higher spending, a sensitive data class, a broader action or dependence that alters the company's ability to operate without the system. The company chooses the specific thresholds in context.

The owner also needs authority to pause. A team responsible for an outcome but unable to limit an unsafe or uneconomic process is accountable in name only. Define who can suspend a workflow, who can restore it and which decisions require escalation. Emergency intervention and permanent retirement need not have the same approval path. The response should reflect the urgency and reversibility of the action.

Delegation should include an evidence obligation. The owner does not have to send every log to the CEO, but must maintain the current result, unresolved assumptions, material exceptions and next decision. The CEO should be able to request the supporting record when an issue becomes material. This separates routine operational reporting from the ability to investigate a consequential claim.

The original approval should expire when a material premise fails, even if the calendar review date has not arrived. A benefit that depended on a certain volume, a control that relied on a particular connector or a service commitment that assumed human review may no longer apply after a change. The owner should know whether the response is to narrow scope, pause expansion or return to the executive decision-maker.

This approach creates room for faster local work. Teams know which changes they can make and which evidence they must preserve. Executives receive fewer routine requests and more focused decisions. The process fails if every minor adjustment becomes material or if materiality is left entirely to the sponsor's convenience. The CEO should review the quality of escalations periodically and refine the boundary based on actual operating experience.

4. Use expectation differences to inspect assumptions

The November keyword reanalysis finds that 70 of 102 historical CEO/founder-classified respondents selected measurable return below six months, compared with 30 of 71 finance-classified respondents. The unrounded difference is about 26.4 percentage points. It is a difference between two selective title groups under a new classification, not a measured change in CEO optimism or a matched disagreement with their own CFOs. [PAY_CEO/founder; PAY_Finance]

The question asks when measurable return is expected, which is not necessarily the date of full investment recovery. The result should prompt a question about assumptions rather than a contest between functions. CEO/founder respondents may work in different company sizes or stages, pursue different projects, define measurable return differently or have different information about implementation. The data does not isolate these explanations. It would be misleading to conclude that one group is more realistic or that a psychological trait explains the difference.

Inside the company, the CEO and CFO can make the comparison useful by examining one investment case together. Ask when the payback clock begins, which costs are included, what benefit category is being promised and what operating action captures it. A forecast can appear to disagree with finance because one person is discussing capacity and another cash, or because one excludes implementation time. Resolving definitions may matter more than negotiating a single preferred number.

The same review should preserve strategic purpose. Some investments create an option to enter a market or improve a service in a way that cannot yet support a precise short-term forecast. The CEO can authorize a bounded learning commitment for that purpose. It should not be relabeled as a near-term cost reduction merely to pass a financial template. The CFO, in turn, can make uncertainty explicit without requiring every experiment to resemble an established production investment.

The November classification also illustrates why executive dashboards need version discipline. A changed grouping rule can alter a percentage even when the underlying answers are unchanged. If the organization changes its definition of an AI initiative, an active user or an accepted benefit, it should record the change and avoid describing the resulting difference as performance movement. Definitions are part of the evidence.

The CEO's role is to insist on a shared decision basis. The goal is not unanimous enthusiasm. It is agreement about the investment's purpose, the evidence available, the assumptions still open and the action

that follows if those assumptions fail. Differences in judgment can remain, but they should be visible enough for the accountable executive to make a deliberate choice.

Exhibit 2. Reclassified historical return-timing answers

Selection	CEO/founder count / base	Share	Finance count / base	Share
Under six months	70 / 102	69%	30 / 71	42%
Six to twelve months	24 / 102	24%	26 / 71	37%
Twelve to twenty-four months	3 / 102	3%	9 / 71	13%
Not sure	6 / 102	6%	8 / 71	11%

Exhibit note: PAY_CEO/founder and PAY_Finance; November keyword classifier v1 applied to historical answers. Multi-select. Classification differs from October; no trend is valid. Groups are not paired executives from the same companies. Small option counts are distribution cells, not separate subgroup estimates.

5. Sequence the portfolio around shared constraints

An AI portfolio can look diversified while depending on one fragile capability. Several applications may rely on the same source data, a small group of reviewers or one integration service. If that dependency fails or becomes overloaded, the apparent spread of projects provides little protection. The CEO should ask which resources and decisions connect the initiatives, not just how many business units sponsor them.

The separate operator instrument identifies data access and quality as the most frequently selected scaling bottleneck, with 29 of 75 answers. Compute appears in 22, integration in 21 and inference cost in 19. These are multi-select operator perceptions, not the constraints of the title-classified CEO group. They provide a reason to examine shared dependencies without proving which dependency limits a particular company. [BOTTLENECK]

A sequencing review should identify the next scarce resource required by each initiative. If two projects need the same subject experts, the organization should decide whether to stagger them, narrow one or add qualified capacity. If several depend on a common data definition, resolving that definition may be more valuable than launching another pilot. The order should reflect what enables useful work, rather than the order in which sponsors obtained funding.

The review should also distinguish reversible experiments from commitments that create dependence. A limited test can generate information with modest exposure. A production rollout that changes staffing, customer promises or a critical process may be harder to reverse. The CEO can favor early experiments that reduce uncertainty before making those commitments, while recognizing that some strategic opportunities require action before every uncertainty is resolved.

An explicit stop or defer decision is part of allocation. If all projects remain nominally active, teams continue to consume attention through updates, requests and partial work. Deferral should state what evidence or dependency would reopen the decision. Retirement should preserve the useful learning and close the associated obligations. This makes the portfolio smaller in practice, not just on a presentation slide.

The CEO should review the opportunity cost in business terms. What customer problem remains unresolved while this initiative uses scarce capacity? Which existing service is exposed when its experts are diverted? Which strategic option becomes easier if a shared constraint is removed? These questions prevent the portfolio from becoming a competition among isolated ROI estimates and make the allocation of executive attention more closely reflect the company's actual priorities.

6. Keep workforce commitments credible

An operating change affects what people are expected to do, how their work is evaluated and what skills they need. The CEO should make those implications explicit rather than allow them to emerge through contradictory messages from different functions. A team told that AI will create learning time may respond differently if every improvement immediately becomes an unplanned workload increase. The company needs a credible account of the change it is asking people to make.

The relevant management decision is broader than how many hours are saved. A workflow can reduce a routine task while increasing the need for judgment, review or customer contact. Job descriptions, training and performance measures may need to change accordingly. If the company continues to reward only the old production count, people may have little reason to invest in the new responsibilities the strategy assumes.

Managers should be able to explain which decisions are settled and which remain open. If a staffing change is being considered, the organization should not make an incompatible promise merely to encourage adoption. If the purpose is to improve service or support growth, the team should see the corresponding priorities and resources. Credibility depends on the operating choices that follow the announcement.

The CEO also needs a channel for evidence that challenges the plan. Employees may discover hidden review work, unsuitable tasks or customer reactions that are absent from the initial business case. A reporting process focused only on success stories can suppress precisely the information needed for responsible expansion. Managers should be expected to surface those findings and show how they changed the workflow.

This does not mean every local concern determines the strategy. Executives still make difficult allocation and organizational decisions. The requirement is to distinguish a considered tradeoff from a benefit assumption that has never been tested. The workforce conversation becomes more useful when it is connected to actual tasks, service needs and evidence, rather than a general prediction that all jobs will change in the same way.

At the executive level, record the commitments that matter: the intended use of released capacity, the training or review support required, the performance measures that will change and the owner of the transition. Revisit them when the workflow expands. The CEO can then assess whether the operating model

is producing the capability the strategy needs, rather than relying on adoption statistics as a substitute for organizational change.

External research and the executive decision

Stanford HAI's October 2026 account of Karunakaran and colleagues' research describes interviews with 184 employees at three firms and a comparison of two law-firm divisions. The more engaged division received both a job-enrichment framing and practical support, including training and exploration time. [E-STANFORD] The comparison should not be reduced to a claim that changing a slogan alone causes adoption. For the CEO, the useful issue is whether the operating support matches the promise made to employees.

The published QJE customer-support study shows that a specific workplace deployment can produce a measurable operating improvement. [E-QJE] Its relevance here is to the level at which evidence is established. A result measured in a particular task and organization should inform a hypothesis, while the company's own evidence determines whether the same operating and financial conditions apply. The CEO should neither dismiss task-level evidence nor treat it as a universal strategy.

OpenAI's work-use research describes people applying AI across occupational task boundaries. [E-OPENAI-WORK] Such usage can create opportunities for broader work, but it does not demonstrate that every resulting output meets a specialist standard. The CEO's delegation decision therefore needs to consider who can judge the result and whether the organization has the capacity to support that judgment. Tool access can widen activity faster than management roles adjust.

NIST's voluntary AI Risk Management Framework places governance alongside mapping, measurement and management. [E-NIST] This provides a useful organizational perspective for the report's contingent-delegation method. An executive should know who owns the decision and which evidence informs it, while technical and operational teams maintain the details within their competence. The framework does not require the CEO to approve every change or validate the OFF sample's findings.

These sources do different analytical work. The workplace research concerns employee engagement and support. The field study concerns measured task performance. The usage research concerns the distribution of activity. The framework concerns organizational risk management. Their combination does not establish a single causal chain from adoption to enterprise profit. It helps identify the separate decisions and evidence needed along that path.

Waiting for mature evidence can mean missing an opportunity. That is why the report distinguishes learning commitments from production commitments. The CEO can deliberately buy information, authorize a limited experiment or preserve a strategic option. The discipline is to state that purpose, bound the exposure and decide what the next result will change. Uncertainty is a condition to manage, not a reason to describe an untested promise as an established benefit.

Source	Executive question informed	Comparison limit
Stanford HAI, October 2026	Does managerial support match the workforce promise?	Framing and practical support are bundled
QJE workplace field study, 2025	What operating result has actually been measured?	Specific task and organizational setting
OpenAI work-use research, July 2026	Who judges work beyond familiar task boundaries?	Usage is not verified competence or ROI
NIST AI RMF	Who owns the decision and its evidence?	Voluntary framework, not survey validation

Working application: decide the sequence, then delegate

Consider an illustrative company with two proposed initiatives. One would automate part of customer onboarding; the other would improve internal account research. Both need the same small integration team and access to a customer-data source whose definitions are inconsistent. The onboarding project has a stronger revenue narrative, while the research workflow is easier to test. These are invented circumstances used to explain the decision method, not an OFF case study.

A simple ranking of projected returns may favor onboarding immediately. The dependency review asks a different question: which initial commitment creates the evidence and capability needed for a responsible next decision? The CEO might authorize a bounded data-definition effort and a narrow research pilot first, while requiring the onboarding sponsor to test its customer assumptions separately. This sequence can reduce uncertainty without pretending that the smaller pilot has proved the larger project's revenue case.

The CEO should also name what will not happen during that stage. The research pilot may remain read-only, serve a limited internal group and carry a capped budget. The onboarding team may conduct customer research without making a new service promise. Those boundaries preserve the opportunity while preventing the early approval from silently authorizing a broader commitment.

Working tool: decision-rights matrix

Decision	Accountable owner	Required input	Escalation condition
Start bounded learning	Business sponsor within delegated budget	Learning question, scope, expiry	Cross-functional conflict or material commitment
Release to production	Operating owner	Accepted outcome, controls, support and cost evidence	Unresolved material risk or dependence
Expand volume / users	Operating owner within agreed boundary	Evidence remains valid at proposed scale	New exposure, cost or service obligation
Expand actions / data	Named executive under company policy	Security, legal and business assessment	Boundary exceeds delegated authority
Pause	Designated operating / incident owner	Trigger and immediate consequence	Material business impact or conflicting obligations
Restore / retire	Authorized business executive	Remediation or closure evidence	Strategic or material resource consequence

The operating owner then maintains a concise review record. It should link to functional evidence rather than copy every detail. Record the workflow purpose, accepted output, current scope, benefit status, material dependencies, unresolved exceptions and next decision. Include the person who can act if the evidence changes. The CEO should be able to use this record to identify the decision requiring attention within a few minutes, while supporting documents remain available for challenge.

Working tool: CEO review and agent register

Field	Entry required	Review responsibility
Business purpose	Customer or operating decision served	Business owner
Current commitment	Budget, users, data, actions and service promise	Operating owner
Evidence state	Observed result, limits and unresolved assumptions	Relevant functional owners
Shared dependency	Scarce people, systems or decisions required	Portfolio / operations lead
Delegated rights	Launch, expand, pause, restore and retire owners	Authorizing executive
Return trigger	Change that requires a new decision	Owner and executive
Next decision	Options, recommendation, date and consequence of delay	Sponsor

At the next review, the CEO should examine whether the pilot resolved the intended uncertainty. If data definitions remain disputed, scaling the onboarding application may still be premature even if the research team likes the tool. If the pilot produces useful evidence and the customer proposition is supported, the next commitment can expand. If the original premise weakens, the company can stop without treating the experiment as wasted effort.

Exhibit 3. Recommended boundaries for executive attention

Situation	Default treatment	Why
Improvement within tested scope	Delegated operating decision	Preserves local speed
Evidence no longer supports a benefit	Reassess funding and claim	Original promise has changed
New consequential action or data use	Renew relevant authorization	Exposure changes even without higher spend
Competition for scarce shared capacity	Executive sequencing decision	Local approvals cannot resolve portfolio tradeoff
Material dependence or customer commitment	Executive review under company policy	Reversal and consequence extend beyond the team

The method succeeds when the CEO makes fewer generic status inquiries and more explicit allocation decisions. It fails when the register becomes a reporting ritual with no effect on funding, sequence or authority. Review the process itself: which decisions were resolved, which were delayed and which escalations could have remained delegated? Adjusting those boundaries is part of the operating model, not an admission that the original design was useless.

The CEO's next agenda

Select the commitments that consume the most scarce management capacity or create the most consequential dependence. Ask each sponsor for the next decision, the evidence supporting it and the shared resource it requires. Resolve conflicts in sequence and ownership before approving another layer of initiatives that compete for the same people.

Then clarify delegation. The operating owner should know what can change locally, what requires renewed authorization and who can pause the workflow. Require a short record of the original promise and the conditions that would invalidate it. The objective is to make decisions faster within a known boundary, while retaining deliberate review of material changes.

Finally, reconcile the business and workforce commitments. The CEO and CFO should agree on the benefit category and financial assumptions; the operating leader should explain the work and capability needed to deliver it. Employees should receive a credible account of what changes in their responsibilities and how the organization will support that change.

The November evidence does not establish a new CEO sentiment trend. It supports a more useful executive conversation: which commitments are justified, what must be learned next and who can act when circumstances change? A CEO who can answer those questions has a basis for allocating capital and attention that extends beyond approving another tool.

Reader questions

Is this a new survey of CEOs?

No. It uses carried-forward common instruments, separate operator context and a disclosed reanalysis of job titles. CEO event rosters are not treated as research answers.

Does the 49% finding measure final production authority?

No. It records CEO mention in a purchase sign-off question, 227/467, multi-select. [SIGN-01]

Did the CEO/finance expectation difference widen?

No valid trend is established. November uses a different keyword classification from October. The 70/102 and 30/71 results are recalculated historical groups. [PAY_CEO/founder; PAY_Finance]

Should the CEO approve every agent change?

This report recommends delegated changes within an agreed boundary, with escalation for material changes to commitments, exposure or evidence.

What is contingent delegation?

It is the report's recommended practice of assigning decision rights that remain valid only while their documented scope and assumptions hold. It is not a survey category or a legal doctrine.

Can an experiment proceed without a reliable ROI estimate?

Yes, as a deliberately bounded learning commitment with a question, cost limit, expiry and next decision. It should not be reported as a verified production return.

What should a CEO do with conflicting functional advice?

Clarify the decision, separate facts from assumptions, identify the tradeoff and assign an accountable choice. Agreement on every judgment is not necessary; a shared evidence basis is.

Citable facts

- CEO mention in purchase sign-off: 227/467, 49%, carried-forward multi-select common instrument. [SIGN-01]

- Historical measurable return below six months under November title classifier v1: CEO/founder 70/102, 69%; finance 30/71, 42%. These are not monthly changes or paired-company comparisons.

- Data access and quality is selected as a scaling bottleneck by 29/75 operator respondents, 39%, separate carried-forward instrument. [BOTTLENECK-01]

Methods and classification note

Calculation-v1 excludes byte-identical exports and invited-only records, then selects the latest eligible nonblank answer per normalized respondent identifier per question within the stated cutoff. Registration timestamps do not verify answer-edit dates. Multi-select answers use any mention; percentages use half-up rounding.

The November classifier examines the first available value among What is your job title?, Job title and Title?. Case-insensitive keyword precedence is CEO/chief executive/co-founder/founder, then CFO/chief financial/finance/controller/treasurer, then technology terms, then investor terms; remaining records are other/unclassified. It does not independently verify current employment or seniority. Mixed titles follow the precedence rule. The script and aggregate group counts are supplied in the collection.

PAY group bases are CEO/founder 102 and finance 71; the remaining 216 respondents are other title groups. The difference below six months is calculated from unrounded fractions: 70/102 minus 30/71, approximately 26.37 percentage points. No significance or causal test is claimed. October used different classification bases, so the November values are a new historical reanalysis, not a continuation of that role series. This version difference is retained in the change ledger.

The selective sample is not representative of all CEOs, CFOs or enterprises. Operator context is not attributed to the title groups. No private executive quotations are used. Management frameworks and the worked company scenario are recommendations and illustrations, not observed respondent behavior.

References

[OFF-DATA] Open Future Forum. Executive AI Statistics, November 2026; title-classification dictionary, aggregate data and change ledger.

[E-STANFORD] Walsh, D. Want Employees to Embrace AI? Stop Selling It as a Productivity Tool. Stanford HAI, 7 October 2026. https://hai.stanford.edu/news/want-employees-to-embrace-ai-stop-selling-it-as-a-productivity-tool

[E-QJE] Brynjolfsson, E., Li, D., and Raymond, L. R. Generative AI at Work. Quarterly Journal of Economics 140(2), 2025. https://academic.oup.com/qje/article/140/2/889/7990658

[E-OPENAI-WORK] OpenAI. How AI is expanding what people do at work. 27 July 2026. https://openai.com/index/how-ai-is-expanding-what-people-do-at-work/

[E-NIST] NIST. AI Risk Management Framework 1.0, official resource page consulted 2 November 2026. https://airc.nist.gov/airmf-resources/airmf/

[OFF-CEO-OCT] Open Future Forum. CEO AI Leverage Report, October 2026, Edition 4. https://openfutureforum.com/research/ceo-ai-leverage-report-october-2026

[MN-TIME] Newlands, M. What Will You Do With the Time AI Gives Back? Murray's Newsletter, 1 November 2026. Included as prior-publication context, not evidence inside the October cutoff. https://murraynewlands.substack.com/p/what-will-you-do-with-the-time-ai

Related research and citation

Read the November CFO report for benefit qualification, AI Leaders for production evidence and Board Preview for materiality and oversight. The CEO report owns cross-functional allocation and delegated decision rights.

Suggested citation: Newlands, Murray. CEO AI Leverage Report, November 2026, Edition 5. Open Future Forum. Evidence cutoff 31 October 2026.

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

Editable working tool

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