Open Future Forum

OPEN FUTURE FORUM

# Executive AI Leverage Report

## NOVEMBER 2026

November 2026 | Edition 5 | Cross-functional research report | Murray Newlands | Open Future Forum

Prepared 2 November 2026. Evidence cutoff: 31 October 2026, 23:59 America/Los_Angeles. Updated security, marketing and investor samples are distinguished from carried-forward buying, founder and operator measures. This edition examines the evidence that must travel between functions before an AI commitment expands.

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/

6,295 substantive body words | 8,687 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, coverage and what the sample can answer
05 1. Make the commitment explicit before allocating capital
06 2. Interpret the new cohorts without manufacturing a common trend
07 3. Preserve the meaning of value as evidence crosses functions
08 4. Decide where released capacity should go
09 5. Renew production evidence when the workload changes
10 6. Connect delegated authority to intervention and service continuity
11 7. Reconcile the clocks on which functions make decisions
12 8. Use one operating record to resolve incompatible claims
13 External research: what supports the thesis and what complicates it
14 A worked executive expansion decision
15 The next executive agenda
16 Executive questions
17 Citable facts
18 Methods and limitations
19 References
20 Related research and citation
21 Editable working tool

## The decision in brief

Expand an AI commitment when the business promise, financial interpretation, operating authority and evidence of accepted results agree on the same scope. Give each handoff an accountable owner and a condition for acceptance. OFF's November evidence shows several kinds of reported impact and concern; it does not prove that those separate observations already form a reliable enterprise operating system.

## Executive summary

An enterprise can have a successful pilot, an approved budget, a signed contract and a security review while still lacking a coherent decision to expand. Each function may have examined a different version of the workflow. The business may expect faster service, finance may expect lower spending, technology may have tested a narrow workload and legal may have approved drafting rather than external action. The missing work is to reconcile those commitments before the organization depends on them.

November adds substantial data in three lanes. The security base grows to 207, with 142 respondents selecting agent identity, access and permissions. The marketing impact base grows to 314, with 141 selecting additional capacity and 141 faster content creation. Investor ownership and impact bases grow to 373 and 364 respectively; 190 impact respondents select better products, 141 cost reduction and 133 help for customers. These are separate selective populations and multiple-selection questions. They cannot be joined into a company-level result. [SEC; MKT; INVOWN; INVIMPACT]

The carried-forward measures remain useful. Proving ROI is selected as a spending blocker by 211 of 389 respondents. Sixty of 75 operator respondents report at least one production agent. In the genuinely matched group of 60 production operators, 30 report real-time visibility into AI running costs, 23 partly and seven no. That matched analysis was already available in October. It is retained as context, not presented as a new discovery or proof of complete workflow cost attribution. [BLOCK; FLEET_AT_LEAST_ONE; OPERATOR_COST]

The report's central recommendation is a common operating record that links existing functional evidence rather than replacing it. The record should identify the business promise, eligible workload, accepted outcome, financial benefit category, authorized actions, operating dependencies and decision rights. Each function remains responsible for its own judgment. The executive team resolves the points at which those judgments describe incompatible commitments.

November's practical contribution is a handoff and decision-clock analysis, followed by an illustrative expansion review for a distributor's quotation workflow. The example shows why released capacity, avoided expenditure and additional commercial contribution must remain separate scenarios until the company has evidence and an operating plan for conversion.

A shared record can become bureaucracy if it repeats reporting without resolving a decision. The recommended record is therefore an index to decision-relevant evidence, with review triggered by changed scope or consequence. It should help the team identify what can proceed, what remains conditional and what would cause the decision to change. More reporting is not the objective; a clearer commitment is.

## What November adds

The November collection reconciles 19 published question distributions to an October-compatible baseline reconstructed from the supplied exports. It adds the later security, marketing and investor sources with respondent overlap resolved separately for each question. The result is a larger cumulative evidence base in those lanes. It is not a new survey of one common executive population, nor a panel tracking the same companies through October.

The first analytical addition is a distinction between an evidence update and an operating trend. The later source can change a cumulative percentage because different people answered, because overlapping respondents supplied a later answer or because the mix of events changed. The report preserves the source-specific results and the cumulative calculation so readers can see what changed without assuming that the enterprise market moved by the same amount.

The second addition is an account of evidence handoffs. A finding useful to one function can lose its meaning when another function receives only the headline. Time saved can become a cash claim, a successful test can become a guarantee of reliability, and a purchase approval can become assumed authority to act. The report identifies the additional evidence required at those transitions and the person responsible for accepting it.

The third addition is a decision-clock method. A financial review, software release, customer commitment and incident response can operate on different schedules. The executive team needs to know which change invalidates an earlier assumption before the next routine meeting. November's proposed record gives each material condition a review trigger and a route for interim decisions.

The fourth addition is a worked expansion decision that preserves alternative uses of capacity. It makes the financial and operating consequences visible without pretending that survey responses establish the result. The example is original to this edition and is clearly separated from empirical findings.

The supplied OFF article on the Stanford research raises the important question of why individual productivity gains do not automatically become company returns. This report uses that question as context and develops the organizational decisions needed to bridge it. It does not reproduce the article's narrative or treat its older OFF percentages as November data. [OFF-STANFORD]

## Evidence, coverage and what the sample can answer

The authorized source package contains 69 exports downloaded on 1 November. Eleven byte-identical copies are excluded, leaving 58 retained exports, 56 nonempty. The retained package contains 9,501 eligible non-invited registration rows and 5,955 unique eligible normalized respondent identifiers. Those network counts describe the supplied file coverage. They are not survey denominators, confirmed attendance or a directly comparable growth measure against October's differently scoped package.

Each question uses the latest eligible nonblank answer per normalized respondent identifier through the cutoff. Question bases therefore differ. A respondent can contribute to several instruments, while many registration rows contain no answer to a given research question. The registry reports exact wording, counts,

bases, multiple-selection status and registration windows. Registration timestamps do not establish when an answer was edited.

The study is selective, self-reported and cross-sectional. It does not estimate representative enterprise adoption, causal effects, audited ROI or a probability-sample margin of error. Broad investor impressions are not company-level results. Planned spending is not actual expenditure. Expected measurable return is not full investment payback. Access-method labels do not establish effective permissions or legal compliance.

The report uses external research to complicate and inform interpretation, with its own populations and limits preserved. Proposed tools are recommendations; numerical scenarios are illustrations. Uncoded private free text is not used to manufacture theme frequencies, budget averages or attributable quotations.

### Exhibit 1. Where the evidence changed

| Instrument | Published October-compatible base | November cumulative base | Treatment |
|---|---|---|---|
| Security priorities and funding | 151 each | 207 each | Later source adds 65 answers; nine overlapping respondents |
| Marketing impact | 218 | 314 | Later source adds 118 answers; 22 overlapping respondents |
| Investor buying ownership | 245 | 373 | Later source adds 151 answers; 23 overlapping respondents |
| Investor impact | 237 | 364 | Later source adds 151 answers; 24 overlapping respondents |
| Common buying, founder and operator questions | Question-specific | Relevant bases retained | Carried forward; no new monthly observation implied |

Exhibit note: Source arithmetic describes unique-response reconciliation, not independent monthly samples. Category counts can also change when an overlapping respondent's later eligible answer replaces an earlier one. Full dates and mappings are in the registry.

# 1. Make the commitment explicit before allocating capital

An AI budget request should identify what the organization is committing to learn or deliver. Funding an experiment, providing an assisted service and making an operating model dependent on automation are different decisions. They can require different evidence and different review schedules. A single approval process that treats all three as equivalent can either block useful learning or permit a consequential commitment before its assumptions are understood.

The common buying data makes financial qualification a relevant question. Proving ROI is selected by 211 of 389 respondents as a blocker. Net-new money is selected as a funding source by 159 of 389, while 120

select uncertainty or no AI budget yet. The questions permit multiple selections. These results show reported concerns and funding approaches within the sample; they do not establish that every initiative should meet the same financial threshold or that lack of an AI-labelled budget means no relevant spending occurs. [BLOCK; FUND]

The executive team should state the investment purpose before choosing the gate. An experiment can justify a bounded cost to answer a consequential question. A cost-reduction program should identify the expense expected to change and the work required to make that change. A customer-service investment may be judged on an accepted outcome and a credible commercial mechanism. A risk-reduction investment may require a different account of exposure and consequence. The evidence should match the decision rather than force every purpose into an immediate cash-saving claim.

Purchase authority is one part of that commitment. In the separate signer question, 227 of 467 respondents select the CEO, 144 finance and 60 no single owner. Those are answers to buying ownership, not verified delegations for production actions. The person who approves the expenditure may not be the person authorized to accept a customer promise, broaden data access or suspend an operating service. The executive record should name those decisions separately. [SIGN]

Expected return timing also needs its original meaning. In the common instrument, 208 of 389 select measurable return within six months. That does not mean the entire investment must have been recovered in cash within six months. A team may observe a useful service result before recovering implementation cost, or observe time savings that never become expenditure reduction. Finance and the business owner should agree what result would satisfy the investment purpose and what further evidence is required for expansion. [PAY]

The CEO's contribution is to resolve competing commitments. If the business case assumes reduced review staffing while security and legal rely on that review as a control, the assumptions cannot all be accepted without further work. The team must redesign the workflow, retain the cost, narrow the promise or accept a different result through the appropriate decision. A budget approval should not leave that conflict hidden for operating teams to discover later.

The practical output is a commitment statement with an owner, an evidence boundary and an explicit next decision. It should distinguish what is approved now from what remains conditional. This allows the company to fund learning and useful deployment while keeping consequential dependencies visible. Capital allocation becomes more precise because the team knows which uncertainty the next expenditure is intended to resolve.

### Exhibit 2. Buying evidence that must retain its meaning

| Measure | Count and base | What it does not establish |
|---|---|---|
| Proving ROI selected as blocker | 211/389; 54% | Audited failure to generate a return |
| Net-new money selected | 159/389; 41% | A market-wide budget share |
| No AI budget yet or uncertainty | 120/389; 31% | Absence of all AI-related expenditure |
| CEO selected as signer | 227/467; 49% | Sole authority over every production decision |
| Measurable return within six months | 208/389; 53% | Full investment payback within six months |

Exhibit note: Carried-forward FUND, BLOCK, SIGN and PAY; multiple selections. Different questions retain different bases and must not be assembled into a matched purchase funnel.

## 2. Interpret the new cohorts without manufacturing a common trend

The new security, marketing and investor sources add useful observations, but they answer different questions. Security respondents identify concerns and funding routes. Marketing respondents identify where AI makes a difference. Investors describe impressions across their portfolios. An executive synthesis should examine how those perspectives inform a decision without implying that they are measurements of the same companies or the same underlying outcome.

Security remains strongly concerned with agent identity and access: 142 of 207 respondents select that category. Dedicated security funding is selected by 74 of 207. Subtracting the two percentages would not measure underfunding. A priority and a funding route are different constructs, respondents can select more than one route, and the survey does not measure the adequacy of the amount spent. The useful question is how the organization funds the controls required by its actual deployment. [SEC; SECFUND]

Marketing's updated impact distribution records 141 of 314 selecting additional capacity, 141 faster content and 130 better customer understanding. Those categories describe reported effects, not a verified increase in accepted output, customer conversion or revenue. Six respondents also select nothing measurable alongside a positive category. The directly calculated positive union is 288 of 314, or 92%. The mixed selections are retained because the instrument does not establish whether respondents mean different teams, different periods or different standards of measurement. [MKT; MKT_POSITIVE; MKT_CONTRADICTORY]

Investor impact has a similar interpretive boundary. Better products is selected by 190 of 364, cost reduction by 141 and help for customers by 133. The positive union is 321 of 364, or 88%, with eight positive-plus-none combinations. The unit is the respondent's view across a portfolio. The result cannot be translated into the percentage of portfolio companies earning a return, because company overlap and

company-level outcomes are not observed. [INVIMPACT; INVIMPACT_POSITIVE; INVIMPACT_CONTRADICTORY]

The later sources also have different registration windows. The security source spans late September through October, marketing begins in September and the investor source begins in August. Calling all their responses October observations would obscure that coverage. The registry separately identifies October registration subsets, but those subsets still do not establish when an answer was edited or create a representative monthly panel.

The executive implication is a set of questions, not a composite index. Are the reported capacity gains producing accepted business outcomes? Does the organization know which actions its agents can take? Can a claimed cost effect be reproduced in a company record? Each question requires evidence from the relevant workflow. A synthesis is useful when it explains those connections while preserving the limits of the underlying measures.

The absence of a common trend is itself important for decision-making. Management should not infer that strong reported impact makes control work less necessary, or that strong security concern means a useful deployment should stop. The observations can coexist. The company needs to evaluate its own commitment using the appropriate outcome, financial and authority evidence, with the new cohort results helping it identify where to ask more precise questions.

### Exhibit 3. Updated results are different kinds of evidence

| Lane | Updated observation | Decision it can inform |
|---|---|---|
| Security | Agent access priority 142/207, 69%; dedicated funding 74/207, 36% | Ask how required controls are funded; do not calculate an underfunding gap |
| Marketing | Capacity 141/314, 45%; speed 141/314, 45%; customer understanding 130/314, 41% | Define accepted output and the customer consequence |
| Investor impact | Product 190/364, 52%; cost 141/364, 39%; customer help 133/364, 37% | Test the relevant company-level mechanism and supporting record |
| Investor ownership | CEO 185/373, 50%; CIO/CTO 97/373, 26% | Identify the actual decision owner in the company under review |

Exhibit note: Updated cumulative, multiple-selection questions. The rows are not comparable prevalence measures and do not form a pooled executive sample.

# 3. Preserve the meaning of value as evidence crosses functions

An operating result changes meaning when it moves into a financial or commercial decision. A reduction in task time can establish a capacity opportunity. Finance then needs to know whether the capacity changes spending, supports additional work or remains unused. Sales may describe a product capability, while the customer contract requires an accepted outcome. The executive team should identify these transitions explicitly because each requires additional evidence.

The first handoff is from technical output to accepted work. A model may produce an answer quickly, but the relevant employee may reject it or spend additional time checking it. A system may report that an action succeeded while the external record remains unchanged. The business outcome should be defined in the workflow where value is supposed to occur. Technical measures remain useful, but they should not silently substitute for acceptance.

The second handoff is from accepted work to an economic consequence. More completed work can improve service or reduce a backlog without changing expenditure. A company may choose to use the capacity for growth, quality or resilience. Those choices should be made by an accountable business owner and reflected in the investment case. An unused capacity estimate should not become a financial benefit merely because it has been multiplied by an hourly labor rate.

The third handoff is from customer interest to a commercial obligation. A pilot can show that a customer values a result without establishing a scalable delivery model or a recurring payment. The founder pricing data includes usage, subscription, outcome and seat components, with overlapping selections. The contract must define what creates payment and what the supplier must deliver. A purchasing team should understand the same terms from the buyer's perspective, including variable cost and acceptance obligations. [PRICE]

The fourth handoff is from commercial records to financial reporting. Usage, contracted commitments, invoices, recognized revenue and cash collected are distinct measures. The appropriate accounting treatment depends on the contract and applicable framework. The executive team should expect a reconciliation where those differences are material, rather than assume that an attractive usage chart is equivalent to revenue quality. This matters to internal forecasts as well as transactions. [E-IFRS15; management application]

These handoffs can expose a valid disagreement rather than an error. Marketing may care about faster production while finance needs a spending consequence. Technology may show a reliable restricted workflow while the business wants a broader service. The task is to identify what additional evidence or decision connects the claims. Forcing all functions to use one headline metric can conceal the differences that management most needs to resolve.

A common operating record should therefore preserve the original unit and interpretation at each stage. It can show task time, accepted output, financial category and commercial result without adding unlike measures together. The executive decision becomes clearer: which result has been demonstrated, which conversion remains a hypothesis and which commitment can responsibly follow? That is the practical bridge between individual productivity and enterprise value.

# 4. Decide where released capacity should go

The workforce consequence of AI is a management choice constrained by the operating system. Time released from one task may be fragmented, concentrated in a role that is not the bottleneck or dependent on continuing review elsewhere. A company cannot assume that adding individual time savings produces an equivalent reduction in staffing. It needs to understand how the work is scheduled and which constraint governs the service.

The common funding instrument includes 73 of 389 respondents selecting money that would have been spent on headcount. That is a reported funding source, not evidence that 19% of employees were replaced or that a corresponding hiring reduction occurred. The operator maturity instrument separately includes an embedded category tied to cost structure or hiring plans if AI were removed. Those questions make workforce assumptions relevant, but neither measures a causal employment effect. [FUND; MATURITY]

Management should distinguish several destinations for capacity. The company may reduce external spending, avoid a planned hire, improve service quality, address a backlog or pursue additional demand. Each destination needs a different action and evidence. Reducing a supplier invoice requires a commercial or operating change. Avoiding a hire requires a credible workload plan. Growth requires demand and the capacity to fulfill it beyond the task that became faster.

Incentives can affect whether the capacity becomes visible. Employees may reasonably worry that reporting a faster task will lead to an unrealistic workload or a loss of support. Managers may preserve the old process while expecting new productivity, leaving staff to perform both. The company should explain the purpose of the change, the work that will stop and the expectations for quality and review. These are management recommendations, not measured attitudes in the OFF dataset.

The Stanford account supplied through OFF's recent article is useful context because it examines adoption in an organizational setting rather than treating access to a tool as sufficient. The intervention described includes support and time alongside framing. The relevant inference is that management arrangements can affect whether employees use a capability, while the account does not establish a universal financial result. This report therefore asks the company to name the destination of released capacity and test whether the surrounding workflow permits it. [E-STANFORD]

The executive team should also preserve capability needed for exceptions. Removing experienced staff can reduce the organization's ability to review unusual cases, train new employees or recover from disruption. A staffing decision should consider those functions, even if routine task time falls. The appropriate choice may still involve a smaller team, but the business case should include the resources required for the service the company intends to provide.

The final decision should state what changes in the operating plan. Which work is removed, which capacity is redeployed, which expenditure changes and which service promise remains supported? If no action follows, the organization may have a useful task improvement without an enterprise financial result. Recognizing that distinction allows management to improve the plan without denying the value of the underlying capability.

# 5. Renew production evidence when the workload changes

The operator sample records meaningful production use, but its questions do not verify reliability. Of 75 fleet respondents, 60 report at least one production agent and 22 more than twenty. The word agent and the reported fleet count are respondent descriptions under the instrument. They do not provide a tested inventory of autonomous systems, the consequences of their actions or the conditions under which they perform well. [FLEET]

The question relevant to expansion is whether existing evidence covers the proposed change. A workflow tested on routine cases may later receive complex exceptions. A drafting service may gain authority to send messages. A new user population may provide different inputs. The product name can remain unchanged while the operating boundary becomes materially different. The release decision should identify which evidence remains valid and which must be renewed.

Reliability should be assessed at the point where the business result occurs. If a workflow changes an external system, verify the intended state and authorization there. If it provides advice or a draft, assess acceptance under the relevant review standard. Anthropic's engineering guidance on agent evaluations distinguishes task outcome from the agent's own narrative, which supports this practical distinction. It does not certify any particular deployment or replace a company-specific evaluation. [E-ANTHROPIC-EVAL]

Production economics also depends on the changed workload. Longer tasks, more retries and additional review can increase cost per accepted outcome even when a model's unit price falls. The direct cost question records real-time visibility into AI running costs, not a complete financial ledger. The matched production-operator result shows 30 yes, 23 partly and seven no out of 60. That result justifies asking what the company can see; it does not prove the completeness or accuracy of any respondent's allocation. [OPERATOR_COST]

Platform and application responsibilities should be separate. A platform may provide authentication, logging or evaluation tools, while the application team defines permitted actions and accepted results. A vendor statement about the platform should not be treated as evidence that a local workflow has been configured or tested correctly. The expansion record should identify which assurance belongs to the supplier and which the company must produce itself.

The company should examine the failures that matter to the proposed commitment. A low-consequence drafting task and an action affecting customer funds require different judgments. Average success can conceal a rare but consequential failure mode. This does not mean every deployment needs an exhaustive universal benchmark. It means the evaluation should match the task, population and consequence that management is proposing to accept.

Requiring a complete review for every change can impede useful iteration. The response is reusable evidence with explicit boundaries and proportionate review of the changed condition. Routine changes can remain delegated. A material change in workload, authority or consequence should trigger the evidence needed for that change, rather than inherit an old approval indefinitely.

### Exhibit 4. Production use and cost visibility, with a valid matched base

| Measure | Result | Interpretation |
|---|---|---|
| At least one production agent | 60/75; 80% | Carried-forward self-reported fleet use |
| More than twenty production agents | 22/75; 29% | Scale category, not a reliability result |
| Cost visibility, full COST base | Yes 34; partly 30; no 13; n=77 | Real-time running-cost visibility question |
| Cost visibility among matched production operators | Yes 30; partly 23; no 7; n=60 | Valid conditional result, already reported in October |

Exhibit note: FLEET, COST and OPERATOR_COST. The matched group comes from respondents answering both relevant questions and reporting at least one production agent. Do not subtract 80% and 44% as a measured visibility gap.

## 6. Connect delegated authority to intervention and service continuity

Authority should describe what the system can do, not merely how it signs in. Per-agent credentials, shared service accounts and delegated identity can support different architectures, but their labels do not establish effective permissions. The organization needs to know the resources, actions and limits available to the workflow, including downstream services that may operate under separate authority.

The operator access instrument has 71 responses, including 22 not applicable yet. Among the 49 applicable responses, 24 select per-agent credentials and 18 shared service accounts. These small, selective counts can inform questions about attribution and control. They do not establish that one group is compliant and another is not, or that a particular identity method caused an incident. The actual implementation must be examined. [ACCESS_APPLICABLE]

The security team's intervention evidence should reach the business owner in operational terms. What stops when access is revoked? What queued or downstream work can continue? Which customer service becomes unavailable? A control can be technically effective at ending one permission while leaving the company unable to meet a service commitment. The executive decision needs both the security result and the consequence for operations.

The legal review should address the authorized act and relevant terms. Permission to draft does not automatically authorize sending an offer or changing a record. A supplier agreement may allocate responsibilities that the application team must implement. The legal and technical records should be compared when the workflow changes. This is a management recommendation; it is not an inference of legal responsibility from the OFF buying-owner question.

Continuity requires a credible operating alternative. That alternative may be a restricted mode, a human process or a different service. It need not duplicate every capability at full scale. The business should state the service level it can sustain, the resources required and the consequences of a longer disruption.

Removing fallback capacity is a strategic and financial choice that should be visible in the same decision as the expected efficiency gain.

CISA's official agentic-AI announcement emphasizes constrained autonomy, identity and oversight. NIST's voluntary framework helps organize the broader risk process. These sources support structured questions, but neither proves that a company's intervention or continuity arrangement works. The relevant evidence is a scoped demonstration and an honest account of what was not tested. [E-CISA; E-NIST]

The decision-rights record should distinguish launch, expansion, restriction, pause, restoration and permanent withdrawal. Different people may hold those rights under the company's arrangements. The executive team should ensure that urgent intervention does not depend on finding the person who approved the original purchase. It should also ensure that restoration requires the appropriate evidence, rather than treating the absence of a visible incident as sufficient reason to resume the earlier scope.

## 7. Reconcile the clocks on which functions make decisions

An AI commitment can change between scheduled reviews. Finance may examine the benefit monthly, a product team may release weekly, a customer contract may renew annually and an incident may require action within hours. A quarterly executive dashboard cannot by itself ensure that the evidence remains valid through those changes. The organization needs triggers that connect the relevant event to the right decision-maker.

The first clock is operational change. Adding a tool, changing a data source or expanding the user population can alter the evidence boundary. The operating owner should identify whether the change affects accepted outcomes, cost, authority or an external commitment. The review should occur before the changed scope is relied on, with routine changes delegated under a clear policy.

The second clock is financial realization. A task improvement may be observed quickly, while expenditure changes only when a contract is renegotiated or a staffing plan changes. Finance should record when the expected conversion can occur and what action is required. A short-term delay in cash benefit may be consistent with the plan; an assumption that never receives an accountable action is a different problem.

The third clock is commercial obligation. A sales or procurement decision can create a promise before the operating team has demonstrated the required service. The company should connect contract approval to the current production scope and cost assumptions. A new response-time commitment, acceptance rule or usage cap can change the economics even without a software release. The relevant owners should receive that change in time to act.

The fourth clock is assurance expiry. A review may remain technically current while its scope no longer covers the deployment. The record should identify the condition that makes evidence stale, rather than rely only on an annual renewal date. A result can remain useful for the original population while being insufficient for a new one. That distinction allows the company to preserve valid evidence without overstating its reach.

The fifth clock is an event requiring intervention or legal assessment. The company should use its established response arrangements and applicable obligations. A generic AI reporting cadence should not displace an incident or disclosure process. The Board and General Counsel reports address the relevant

distinctions in greater detail; this flagship's concern is the handoff that ensures the right process receives the facts promptly.

The executive team should agree which clock governs the next commitment. If a customer launch precedes completion of a material control test, the schedule conflict needs an explicit decision. If a benefit depends on a contract change next quarter, the forecast should reflect that timing. The operating record should make these dependencies visible so that one function's deadline does not silently override another function's necessary condition.

### Exhibit 5. Decision clocks and review triggers

| Clock | Trigger | Owner of the next decision |
|---|---|---|
| Operating change | New workload, data, tool or action | Business and production owners, with relevant control review |
| Financial conversion | Spending, hiring or revenue action required | Finance and accountable business executive |
| Commercial commitment | New promise, acceptance rule or pricing term | Commercial owner with legal, finance and operations |
| Evidence validity | Deployment exceeds the tested or reviewed boundary | Owner of the affected assurance |
| Incident or material exposure | Facts meet an established escalation condition | Designated incident, legal and executive decision-makers |

Exhibit note: Proposed management framework. No universal review interval, legal deadline or measured respondent practice is asserted.

## 8. Use one operating record to resolve incompatible claims

The common record should be small enough to use and complete enough to expose disagreement. It should link the business promise to the accepted outcome, financial category, authority boundary and operating dependencies. Supporting evidence can remain in the systems where it is maintained. The record's purpose is to identify the current decision and the conditions on which that decision rests.

Each handoff should state what the receiving function accepts. Finance can accept that an observed time reduction is a capacity result while withholding a cash interpretation. Security can accept a restricted access design while requiring a downstream intervention test. Legal can accept a specified drafting use while reserving judgment on external action. These are useful qualified conclusions, not failures to collaborate. The executive team needs them to understand what is actually ready.

When evidence conflicts, the team should identify whether the disagreement concerns facts, definitions or judgment. Different completion definitions can produce different valid counts. A missing cost can change the financial interpretation. Different risk tolerances can produce a decision disagreement even when the facts

are shared. Each type requires a different response. More data will not necessarily resolve a judgment that management must make explicitly.

Unsupported claims should enter a repair process. The company should identify the missing evidence, the decision constrained and the person who can resolve the issue. Closure can support the claim, narrow it, retain it as a hypothesis or withdraw it. The Investor report develops this process after capital is committed; the same discipline is useful inside an executive team when an early claim has acquired more certainty than its evidence supports.

The record should preserve corrections. If a financial benefit or performance claim changes, the team should identify where the earlier version was used and update the relevant materials through the company's normal process. A corrected spreadsheet is insufficient if the old number still supports a staffing decision or customer promise. The consequence of the correction is part of the work, not an optional communication step.

Material decisions should reach the board through the appropriate governance route. The board needs to understand the commitment, the evidence, the alternatives and the unresolved consequences. It does not need every operational event. Management should explain which conditions remain within delegated authority and which changes require further oversight. The Board Preview retains that distinction without treating operator responses as director sentiment.

The test of the record is whether it improves a real decision. Can the team say what may expand now, what remains restricted, who can intervene and what evidence would change the conclusion? If the record only adds fields to an existing dashboard, it should be simplified or redesigned. Its value lies in resolving incompatible commitments before they become embedded in the company's finances, service obligations and operating model.

### Exhibit 6. Evidence handoff and acceptance record

| Handoff | Evidence supplied | Acceptance question |
|---|---|---|
| Product to business owner | Task result, population and exception record | Is this accepted work in the intended workflow? |
| Business owner to finance | Capacity destination and operating action | Which financial category is supported, and when? |
| Security to business owner | Effective permissions and intervention result | What stops, what continues and what service is affected? |
| Technology to counsel | Actual data, actions, recipients and changes | Does the deployed use fit the assessed authority and terms? |
| Commercial team to operations | Customer promise and acceptance conditions | Can the current service deliver the commitment at the modeled cost? |
| Management to board | Material commitment, alternatives and uncertainty | Is the decision within the approved strategy and delegation? |

Exhibit note: Proposed cross-functional tool. Each function retains responsibility for its own judgment; the table is not a certification or a substitute for company governance.

## External research: what supports the thesis and what complicates it

The QJE study Generative AI at Work reports evidence from a specific customer-support setting, with heterogeneous effects across workers. It supports examining where assistance changes performance, while limiting any claim that the same effect applies across jobs or companies. OFF's capacity and impact selections do not reproduce the study's design. The comparison is useful for framing a company-specific test, not importing an effect size into a business case. [E-QJE]

METR's February 2026 update makes measurement uncertainty explicit in its newer developer-productivity work. Selection and changing work patterns complicate a reliable current estimate. This matters because an executive can otherwise treat a familiar headline as a stable property of AI. The appropriate inference is to examine the task, population and method when using productivity evidence, including evidence produced inside the company. [E-METR]

Stanford HAI's October article describes research involving 184 employees at three firms. Its account connects adoption with framing and organizational support; the described intervention bundles support such as training and time. It should not be presented as proof that wording alone causes a particular financial return. The executive implication is to examine the conditions under which employees can use a tool effectively. [E-STANFORD]

The HBS AI Institute's institutional account of the GenAI wall emphasizes the interaction between tasks and expertise. This report relies on that accessible institutional account and does not repeat unverified detailed

estimates from the underlying working-paper PDF. It complicates a simple claim that better tools remove the need for domain knowledge. The relevant management question is where review and expertise remain necessary in the workflow. [E-HBS-WALL]

OpenAI's July 2026 work analysis uses a large set of work-related messages to examine how people use AI across occupational tasks. Message-level use is different from employee productivity, accepted work or company returns. The finding is relevant to changing task boundaries, but the company still needs evidence that the additional work is appropriate and useful. Broader use does not by itself establish competence in a new domain. [E-OPENAI-WORK]

Deloitte's third-quarter CFO Signals article adds a large-company finance perspective on cost governance and intervention. Its population of 200 North American CFOs at companies with at least $1 billion in revenue differs from OFF's samples. The comparison supports asking about control and response capability, while providing no basis for declaring a measured gap between the two surveys. [E-DELOITTE3]

FinOps guidance supplies a distinction between technical and business units in cost analysis. Anthropic's evaluation guidance distinguishes an agent's account from the resulting state. NIST and CISA provide risk and security guidance with their own scope. These sources inform the proposed handoffs, but they do not validate the OFF survey or certify the worked example. Their role is to help specify the evidence a company should obtain. [E-FINOPS; E-ANTHROPIC-EVAL; E-NIST; E-CISA]

Together, the sources support a more conditional interpretation than either enthusiasm or dismissal. AI can improve particular tasks, while the effect depends on the work, people and operating arrangements. A company may also create valuable new capacity without reducing cost immediately. The executive decision is to identify the mechanism it intends to realize and the evidence required at each transition. The next action should follow the uncertainty that matters most to that mechanism.

### External comparison table

| Source and evidence type | Relationship to OFF | Practical implication and limit |
|---|---|---|
| QJE, 2025 field study | Supports task-specific performance inquiry; different design | Test relevant populations; do not import a universal uplift |
| METR, February 2026 research update | Complicates simple productivity headlines | Preserve selection and measurement limits |
| Stanford HAI, October 2026 research account | Adds organizational conditions of adoption | Separate bundled support from a pure framing effect |
| HBS AI Institute, September 2025 institutional account | Complicates claims that tools replace expertise | Identify review and domain-knowledge dependencies |
| OpenAI, July 2026 message analysis | Adds evidence of task-use breadth | Messages are not verified employee or company outcomes |
| Deloitte, October 2026 CFO survey article | Adds a different finance population | Ask comparable questions without pooling the samples |
| FinOps, Anthropic, NIST and CISA guidance | Informs proposed operating methods | Guidance is not an observed OFF result or certification |

Exhibit note: Sources, dates and scope are specified in references. The comparison is analytical; no pooled effect estimate or claim of direct survey replication is made.

## A worked executive expansion decision

Consider a hypothetical industrial distributor processing 6,000 quotations per month. An AI-assisted workflow reduces the assumed average drafting time from eighteen minutes to ten for the eligible workload, with accepted quality held constant for this illustration. The arithmetic releases 48,000 minutes, or 800 staff hours. These figures are invented to demonstrate a decision method. They are not an OFF observation, a vendor result or a forecast.

The business initially describes the 800 hours as a saving. Finance asks which expenditure will change. Operations explains that only 400 hours are usable at the times and locations where additional work exists. The company could remove two temporary assignments of 160 hours each at $35 per hour, if the relevant service obligations remain supported. That would reduce monthly expenditure by $11,200. The remaining capacity is not automatically another cash benefit.

Assume the expanded workflow requires $10,000 per month of AI service and infrastructure plus $4,000 of additional review and support. Under the temporary-spending scenario, the simplified recurring cash effect is $11,200 less $14,000, or negative $2,800. The workflow may still improve service, but the proposed cash-saving claim is not supported by these assumptions. Management should identify the service benefit separately rather than relabel the released hours as cash.

A different scenario uses the available capacity to support additional demand. Suppose management hypothesizes 300 additional fulfilled orders per month with $80 contribution each before the incremental AI costs in this example. That would produce $24,000 of contribution, less $14,000 of AI, review and support cost, leaving $10,000. The demand, conversion and fulfillment assumptions require evidence. This scenario retains the staffing capacity and does not also count the $11,200 temporary-spending reduction.

### Exhibit 7. Illustrative capacity destinations

| Item | Spending-reduction scenario | Additional-demand scenario |
|---|---|---|
| Task-time capacity released | 800 hours | 800 hours |
| Capacity considered usable in the operating plan | 400 hours | 400 hours |
| Financial conversion assumption | Remove 320 temporary hours at $35 | Fulfill 300 additional orders at $80 contribution |
| Gross monthly financial effect | $11,200 | $24,000 |
| Incremental AI, review and support | $14,000 | $14,000 |
| Simplified recurring result | ($2,800) | $10,000 |
| Evidence still required | Actual spending reduction and service continuity | Demand, conversion, fulfillment and contribution |

Exhibit note: Entirely hypothetical; scenarios are alternatives and must not be added. Contribution in the demand scenario is stipulated before the incremental AI costs shown. Taxes, financing, working capital and other effects are excluded.

Now add an operating change. The business wants the agent to send quotations directly rather than prepare drafts for employee approval. Technology's original evaluation covered drafting, while legal's review assumed employee approval before external communication. Security has tested the current application's access but has not yet demonstrated the intervention path for the new sending capability. The financial scenarios do not resolve those conditions.

The executive team should separate the decisions. It can continue the existing drafting scope while testing the commercial hypothesis and evaluating the proposed sending authority. The business owner defines the eligible quotation and acceptance rules. Finance validates the financial conversion and omitted costs. Technology tests the changed workflow. Security demonstrates the relevant intervention. Counsel assesses the actual proposed act and terms. The CEO resolves any conflict between the desired launch date and those conditions.

## Exhibit 8. Decision and ownership matrix for the illustration

| Decision | Accountable lead | Evidence before the decision changes |
|---|---|---|
| Continue restricted drafting | Business owner | Accepted output and current operating limits |
| Count a spending reduction | Finance with operations | Actual expenditure action without unsupported service loss |
| Count additional commercial contribution | Commercial and operations leaders with finance | Fulfilled demand and a supported contribution calculation |
| Permit external sending | Authorized business executive | Relevant technical, security and legal conditions resolved |
| Pause consequential actions | Designated incident or operating authority | Trigger defined in the response arrangement; actual containment scope verified |
| Restore or expand after a problem | Authorized executive with relevant reviewers | Evidence that the condition causing restriction has been addressed |

Exhibit note: Proposed allocation for the hypothetical company. Actual authority must follow the organization's governance and applicable obligations.

If implementation requires a further $60,000, management should keep that one-time amount separate from the recurring scenarios. Under the assumed positive $10,000 monthly contribution, simple recovery would take six months after reaching that steady state. That is a transparent illustrative calculation, not the meaning of the OFF survey's measurable-return question. A later start, lower demand or higher review cost would change the result and should be reflected in the plan.

The common operating record should contain the promise, workload, accepted outcome, financial category, authority boundary, relevant terms, dependencies and next decision. It should link to the supporting evidence and identify the unresolved assumptions. The team does not need a new database for the illustration to work. It needs a shared account of which commitment is currently supported and which remains conditional.

At the next review, the decision should change according to what has been learned. If demand is absent, the commercial scenario weakens. If review cost persists, the economics change. If the sending capability cannot meet the required boundary, drafting may remain useful. A smaller supported use case can be a sound outcome. The executive team's responsibility is to make that judgment explicitly rather than preserve the original expansion story at all costs.

# The next executive agenda

Choose the material AI commitment whose expansion is most imminent. Ask the business owner to state the promised result and eligible workload. Ask finance to identify the supported benefit category and the action required for conversion. Ask technology and security to identify the evidence boundary and intervention capability. Ask counsel and the commercial owner to identify the relevant authority and external obligations. Bring the disagreements into one decision.

The agenda should end with a clear statement of what may proceed, what remains conditional and who can change the scope. Each unresolved condition should identify the evidence required and the decision it constrains. Avoid creating a long action list in which every item is treated as equally important. Focus on the assumption that determines the next consequential commitment.

Use the November data to improve the questions, while preserving its limits. Updated marketing and investor impact can coexist with security concern and incomplete cost visibility. Those observations do not cancel one another or establish a common market score. They point toward the operating work required to turn a useful capability into a result the enterprise can understand and sustain.

The central task is to keep the promise, evidence and authority aligned as the workflow changes. When those elements agree, management can delegate and expand with a clearer basis. When they diverge, the company should qualify the claim, repair the evidence or change the commitment. That is the practical work between a faster task and a defensible enterprise result.

# Executive questions

### Are all November figures new observations?

No. Security, marketing impact and investor instruments are updated cumulatively. Common buying, founder and direct operator measures are carried forward unless specifically labelled as historical reanalysis.

### Can the functional samples be pooled into one enterprise result?

No. They use different questions and populations, with possible overlap. Their bases are not additive unique-company counts.

### Does expected measurable return mean full payback?

No. PAY asks about measurable return timing. A full investment-recovery calculation is a separate financial analysis.

### Does the matched cost result prove complete workflow cost attribution?

No. It concerns self-reported real-time AI running-cost visibility. Thirty of 60 matched production operators answered yes. [OPERATOR_COST]

### Does the difference between access concern and dedicated funding measure underfunding?

No. Priorities and funding routes are different constructs, with multiple selections and no measure of funding adequacy. [SEC; SECFUND]

### Can released capacity and the resulting cash saving both be added as benefits?

Not when they represent the same underlying improvement. The financial conversion must be identified, and alternative uses of the same capacity must remain separate.

### Does this report contain a new representative executive benchmark?

No. It uses selective registration responses, verified external context and labelled recommendations. No probability-sample inference or causal enterprise return is claimed.

### What should trigger a new decision?

A material change in the workload, authority, external commitment, financial assumption or validity of the evidence supporting the current scope.

## Citable facts

- Security agent-access priority: 142/207, 69%, updated cumulative. [SEC]

- Marketing capacity and content speed: 141/314 each, 45%, updated cumulative. [MKT]

- Investor better-product impact: 190/364, 52%, updated cumulative. [INVIMPACT]

- Proving ROI as a blocker: 211/389, 54%, carried forward. [BLOCK]

- Matched production operators reporting yes to real-time AI running-cost visibility: 30/60, 50%; already reported in October. [OPERATOR_COST]

- The supplied file package contains 58 retained exports after 11 byte-identical copies are excluded; network coverage is not comparable to October's differently scoped total. [OFF-DATA]

## Methods and limitations

The canonical calculation-v1.1 uses 69 authorized exports downloaded on 1 November 2026, through 31 October at 23:59 America/Los_Angeles. Eleven byte-identical exports are excluded. Eligible records exclude invited-only rows and include other application statuses; they are not confirmed attendees. The latest eligible nonblank response is selected per normalized respondent identifier per question, with deterministic filename and row tie-breaking. Registration timestamps do not establish answer-edit timestamps.

The 19 mapped published baseline questions reconcile exactly in counts and bases when the added cohorts are excluded from the current snapshots. This is a reconstruction from the supplied snapshots, not proof that historical answer-edit states were preserved. FINSTAGE is an additional historical instrument newly

documented in the registry, not a new November observation. Title-classified subgroup analysis uses a documented classifier and is not comparable to the earlier classifier without qualification.

Multiple-selection results are any mention. Positive unions and contradictory combinations are computed directly. Matched and applicable results use their stated conditional bases. Small categories are not promoted into unsupported comparisons. Percentages use unrounded values for calculations and half-up rounding for display. Different question and source populations are not pooled into a common outcome.

The data is selective, self-reported and cross-sectional. No causal effect, audited return, company-level prevalence, representative market size or statistical significance claim is made. Planned spending, reported impact, expected return, maturity and access method retain their exact conceptual limits. Private respondent identifiers and uncoded sensitive free text are excluded from the publication package.

External sources are assessed within their own dates, populations and designs. Institutional summaries are identified as such where the underlying full paper was not used. Proposed operating tools and the distributor example are original recommendations and illustrations. The illustration's scenarios are alternatives, with explicit cost and timing assumptions; they are not empirical OFF findings.

## References

[OFF-DATA] Open Future Forum. Executive AI Statistics, November 2026, calculation-v1.1. Complete distributions, question dictionary, source manifest, comparison ledger and aggregate CSV/JSON.

[E-QJE] Brynjolfsson, E., Li, D. and Raymond, L. Generative AI at Work. Quarterly Journal of Economics, 140(2), 889-942, 2025. https://academic.oup.com/qje/article/140/2/889/7990658

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

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

[E-HBS-WALL] Harvard Business School AI Institute. Why AI Helps - Until It Doesn't: Inside the GenAI Wall Effect. 18 September 2025. Institutional account used; no unverified detailed estimates from the underlying working-paper PDF are reproduced. https://aiinstitute.hbs.edu/why-ai-helps-until-it-doesnt-inside-the-genai-wall-effect/

[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-DELOITTE3] Deloitte. CFO Signals, third quarter 2026. 21 October 2026. https://www.deloitte.com/us/en/insights/topics/business-strategy-growth/3q-2026-cfo-signals-survey.html

[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-CISA] CISA and partners. Official announcement of Careful Adoption of Agentic Artificial Intelligence Services. 1 May 2026. https://content.govdelivery.com/accounts/USDHSCISA/bulletins/41544ff

[E-IFRS15] IFRS Foundation. IFRS 15 Revenue from Contracts with Customers, official overview consulted 2 November 2026. https://www.ifrs.org/issued-standards/list-of-standards/ifrs-15-revenue-from-contracts-with-customers/

[OFF-STANFORD] Newlands, M. AI Is Making Employees Faster. Why Isn't Every Company Seeing the Returns? Open Future Forum, 31 October 2026. https://openfutureforum.com/blog/ai-productivity-returns-stanford-karunakaran

## Related research and citation

The November collection contains the CFO, CEO, CISO, CMO, AI Transformation, AI Leaders, Board Preview, YC Founder, VC and CVC, Investor, Private Equity, Investment Banking and General Counsel reports, plus the complete Executive AI Statistics registry. Each owns a distinct executive decision; the registry is the source of record for the figures.

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

Open Future Forum research and prior editions: https://openfutureforum.com/research/

## Editable working tool

Executive commitment handoff: download the editable CSV worksheet. The worksheet implements the report's recommendation; it is not a validated score, mandatory standard or completed assessment.