Open Future Forum

OPEN FUTURE FORUM

AI Transformation Report

NOVEMBER 2026

November 2026 | Edition 4 | Research report | Murray Newlands | Open Future Forum

Prepared 2 November 2026. Evidence cutoff: 31 October 2026, 23:59 America/Los_Angeles. The direct maturity instrument is carried forward. This edition develops the operating decisions required before accepting greater dependence on AI; it does not claim a new transition rate.

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/

4,357 substantive body words | 5,597 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 to the baseline
04 The population behind the stage labels
05 1. Decide what must change in the process
06 2. Design a pilot that can change a decision
07 3. Assign the destination of the benefit
08 4. Accept dependence deliberately
09 5. Redesign work without assuming expertise transfers automatically
10 6. Keep the transformation record current after launch
11 External evidence: the boundary between task help and organizational change
12 Applying the gates to a real operating decision
13 The next transformation review
14 Reader questions
15 Citable facts
16 Methods and limitations
17 References
18 Related research and citation
19 Editable working tool

The decision in brief

Transformation should be judged by a changed operating process and a defensible business result, with dependence accepted deliberately. In OFF's direct maturity instrument, 12 of 91 respondents selected embedded AI under a specific cost-structure or hiring-plan definition. November explains how to examine that commitment without treating deployment, autonomy or a higher stage label as a goal in itself. [MATURITY]

Executive summary

The carried-forward maturity distribution contains 19 respondents exploring, 18 piloting, 42 deployed in production and 12 embedded. The embedded option is defined as a condition in which removing AI would change the respondent's cost structure or hiring plan. It is a stronger statement than using a tool, but it remains a self-reported condition rather than a verified assessment of business dependence. [MATURITY]

The combined deployed-and-embedded count is 54 of 91, or 59%. That is a cross-sectional description. It does not show that 59% completed a transformation journey, that the 42 deployed respondents will become embedded, or that one management practice caused their stage. November has no new comparable maturity answers, so these figures retain their historical observation status. [MATURITY_DEPLOYED_OR_EMBEDDED]

The practical question is what the organization should change before relying more heavily on the workflow. A successful task-level pilot can leave the surrounding process unchanged. A faster input may wait in the same approval queue; a useful draft may require more specialist review; a capacity gain may have no assigned destination. The transformation leader must examine those handoffs and the operating decisions needed to capture value.

This edition adds a dependence-acceptance review. It asks what the business can now do, what it has stopped doing, which assumptions support that change and what commitment becomes difficult to reverse. The review connects process design, economics, people and continuity, while leaving technical release decisions and formal benefit qualification to the relevant functional owners.

A staged review can favor narrow efficiency projects over strategic change. The report therefore distinguishes evidence for learning from evidence for dependence. An organization can fund uncertainty deliberately and preserve an option without claiming an established operating return. It should require stronger evidence before changing staffing assumptions, customer promises or critical processes around that option. The aim is a transition the organization can explain and sustain, rather than a maturity label that rises faster than its operating capability.

What November adds to the baseline

October introduced the direct four-stage maturity question. November's export reproduces its 91-person distribution, with no new comparable answers in the supplied records. This is a carried-forward baseline. The report preserves the exact embedded definition and does not replace it with a looser meaning such as daily use, widespread adoption or autonomous execution. Those may be relevant observations, but they are not what this question measured.

The first addition is a process-handoff analysis. The report examines why improving one task can leave the end-to-end service unchanged and how to identify the next constraint. This goes beyond reporting deployment stages: it asks which decision, queue, review or data dependency must change for a local improvement to become a useful operating result.

The second addition is a dependence-acceptance record. October's continuity drill already examines interruption and recovery. November addresses the earlier management choice to make the business dependent: which manual capacity, contractual flexibility or alternative process will be reduced, and what evidence justifies that reduction? A continuity plan is relevant to this decision, but it is not the whole transformation case.

The third addition is a review of the destination of benefits and the credibility of workforce commitments. A transformation program should identify the service, growth, quality or financial result it intends to capture and assign an owner for the operating change. It should also explain the support people need to perform their revised work. These are recommendations, not findings that a particular share of respondents followed such a process.

The stage gates in this report are a management framework. They are kept separate from the survey's four response categories and from any five-stage advisory model used in a market map. The framework does not turn the cross-section into a measured journey or imply that embedded AI is always the right destination. A bounded production workflow may be the appropriate final state if deeper dependence adds little value or creates disproportionate exposure.

The population behind the stage labels

The maturity question comes from the Enterprise AI at Microsoft application flow. The current export contains 314 eligible non-invited application rows, but only 91 eligible unique respondents answered the maturity question. Other operator questions have different completion bases. The report uses the question base for the stage distribution; it does not describe the result as the maturity of all applicants, attendees or companies.

The sample is selective and self-reported. Multiple respondents can represent one organization, and an answer about a function is not automatically a company-wide assessment. The instrument asks where AI is in the respondent's function today. The publication therefore avoids turning a functional stage into a claim that the entire enterprise has transformed.

Answers are selected separately per question by normalized respondent identifier and the latest eligible nonblank registration record within the cutoff. Registration dates are not verified answer-edit dates. The four maturity categories are single selection. The analysis does not observe the same function progressing through stages, the duration of a pilot or a causal relationship between a management practice and a later outcome.

The cost, fleet and bottleneck measures provide separate operator context. Where a matched result is used, its matched base is stated explicitly. No difference between unrelated percentages is called an operating gap. Recommendations below are intended to help an organization produce the evidence that the survey cannot supply about its own processes, benefits and readiness for dependence.

1. Decide what must change in the process

Transformation begins with an operating problem, not a target stage. The leader should identify what the organization cannot currently do well enough: complete a service quickly, understand a customer need, handle volume reliably or make a decision with suitable evidence. The proposed AI workflow should then be evaluated as one way to change that process. A project can use advanced technology while leaving the underlying problem untouched.

Map the process around the proposed task. What arrives before it, what must be decided afterward and where work waits? A drafting improvement may have little effect if approval remains the constraint. A classification improvement may not help if the source data is incomplete. A research assistant may save time while the organization continues to make decisions through a calendar or committee process that does not use the earlier result.

The transformation leader should ask the receiving team what it needs to act. An output that looks complete to the producing team may lack the evidence, format or authority required by the next owner. Redesigning that handoff can be more important than improving the model's isolated task score. The acceptance definition should therefore include the downstream use, not just whether the AI produced a plausible response.

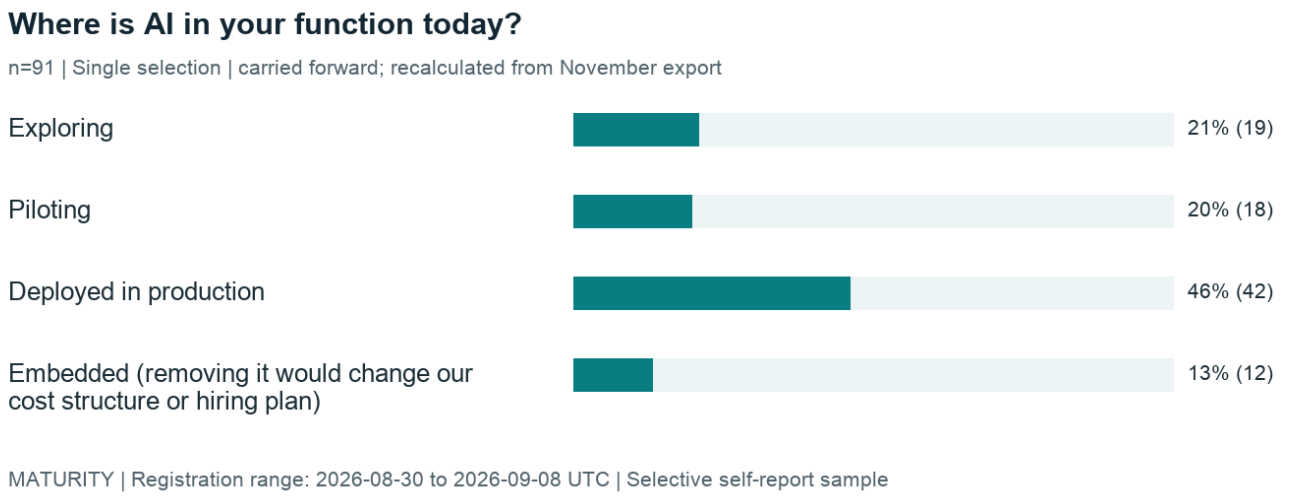
The baseline should include waiting, review and exceptions. Average handling time alone can miss the cases that consume specialist capacity or create customer frustration. The organization should know whether the new workflow changes the distribution of work, including which cases are escalated and whether difficult work is shifted to another team. An apparent improvement at one boundary can be a cost transferred elsewhere.

The next change should target the actual constraint. If the approval queue dominates completion time, the answer may be a clearer decision rule or a better evidence packet. If data quality prevents reliable use, the next investment may be in source ownership. If the workflow produces more output than the receiving team can absorb, expansion may need to wait. The transformation program should be willing to change its plan when the process evidence identifies a different problem.

This is the first gate: demonstrate how the proposed intervention changes the end-to-end process for a defined purpose. The evidence can be modest during exploration, but the causal story should be explicit.

Without it, the organization risks deploying a useful tool and then attributing broader operating improvements to it without knowing whether the relevant handoffs changed.

Exhibit 1. The direct four-stage maturity distribution



Exact response	Count	Base	Share
Exploring	19	91	21%
Piloting	18	91	20%
Deployed in production	42	91	46%
Embedded (removing it would change our cost structure or hiring plan)	12	91	13%

Exhibit note: MATURITY; carried-forward single-selection functional maturity instrument. The distribution is cross-sectional, not a transition history. Embedded is defined by the wording shown, not by tool usage or autonomy alone.

2. Design a pilot that can change a decision

A pilot should resolve an uncertainty that matters to a real decision. The sponsor might need to know whether the workflow meets a quality standard, whether users can operate it with available support or whether a dependency makes expansion uneconomic. Those questions require different observations. A general instruction to try AI can produce enthusiastic demonstrations without answering any of them.

Define the comparison before running the pilot. Record the task population, accepted output, observation period and important exclusions. Where a controlled comparison is feasible, use it. Where it is not, explain what else changed and which inference remains uncertain. A historical comparison may still be useful, but it should not be presented as an experiment that isolated AI's effect.

The pilot should include ordinary and difficult work in a proportion appropriate to the intended use. Selecting only convenient examples can establish that the workflow is possible without showing that it is dependable.

At the same time, a test consisting only of extreme failures may not describe normal operating economics. The leader needs a deliberate test set and a record of what it covers, with separate attention to consequential exceptions.

Participation can also affect the result. Volunteers may have more interest, skill or time than the wider group expected to adopt the workflow. A successful pilot should therefore record the support participants received and the capabilities needed to reproduce the result. Expansion may require training, clearer interfaces or a narrower use case. Those are part of the operating design, not incidental overhead to be removed from the business case.

The decision rule should allow more than success or failure. The evidence may support expansion, another bounded test, a narrower scope, a redesigned process or retirement. A pilot that identifies a costly dependency can be valuable if it prevents an unsuitable commitment. The program should retain the finding rather than reinterpret the original objective so that every pilot appears successful.

Before the pilot begins, name the person who can make the next decision and the resources that decision would require. Otherwise, the team may produce convincing evidence that cannot be acted upon. The transformation leader's role is to connect learning with authority, ensuring that the pilot purchases information the organization is actually prepared to use.

3. Assign the destination of the benefit

A faster task creates an opportunity, but the organization still has to decide what happens next. The released capacity might improve service, absorb growth, reduce overtime, avoid a hire or support new work. Those destinations have different evidence requirements and different owners. A transformation program should identify the intended destination before consolidating the benefit into a portfolio total.

The separate marketing instrument illustrates the distinction without resolving it. Respondents report capacity, faster content and customer understanding as areas of impact. Those descriptions indicate where people perceive change; they do not establish a financial conversion. The transformation leader should use such signals to locate workflows worth examining, then require the organization's own evidence for the benefit it intends to claim. [MKT]

The receiving capacity matters as much as the released capacity. If a team produces more analysis but decision-makers have no time to use it, the business may not gain much. If a service can process more demand but sales or onboarding remains constrained, growth may not follow. A benefit plan should identify the complementary change needed to use the improvement and who has the authority to make it.

Avoid assigning the same improvement to several destinations without a consistent operating model. Hours used to avoid a hire cannot also be assumed to produce additional customer work at the same time. A project may divide the benefit, but the allocation should reconcile to the workload and staffing evidence. The CFO report provides the financial qualification method; the transformation record should show the underlying operational destination.

The benefit may also change over time. Early deployment can require more support and review, while later operation may become more efficient. Conversely, a workflow can become more expensive as difficult

cases, broader users or new requirements enter the scope. The program should preserve the original assumptions and review the result at meaningful operating intervals rather than rely indefinitely on the pilot's best observed period.

The gate is a named, plausible and measurable benefit destination. It does not require every project to promise cash savings. It requires management to explain how the operating change serves the intended purpose and what evidence would show that it did not. This makes the transformation claim more concrete than a general assertion that employees are more productive or that the organization is becoming AI-first.

4. Accept dependence deliberately

The embedded category in the OFF instrument is specific: removing AI would change cost structure or hiring plans. Twelve of 91 respondents selected it. That is evidence of reported dependence under the question's definition, not proof that the dependence is beneficial, resilient or well controlled. A company can become dependent on a weak process as well as a strong one. [MATURITY-04]

The transformation leader should therefore treat dependence as a commitment to evaluate. What capability will the organization reduce or stop maintaining because the AI workflow is expected to remain available? Which staffing assumption, customer promise or service design relies on it? Which supplier or data dependency becomes material? These questions identify the consequences of embedding before they become difficult to reverse.

Dependence can be economically sensible. Maintaining a full duplicate manual process indefinitely may be expensive and may prevent the organization from capturing the benefit it sought. The decision should compare the value of the new operating model with the cost and effectiveness of alternatives. The goal is not to avoid reliance on technology; it is to understand and authorize the reliance the business chooses.

The evidence should include what happens when the workflow is unavailable or unsuitable for a case. October's continuity drill offers a method for testing interruption. November's acceptance review asks what management does with that result. A known fallback that meets only a limited service level may be acceptable for one process and unacceptable for another. The business owner must make that tradeoff explicit rather than assume that any documented fallback is sufficient.

The commitment should also account for change in the supplier or operating environment. A new pricing structure, degraded source data, altered terms or a changed model can affect the basis for dependence. The organization should identify which changes reopen the decision and what flexibility remains. A workflow can be embedded while still preserving an exit or adaptation path; those are operating choices that require resources.

The gate is a documented acceptance of dependence by the appropriate business executive, supported by current economic, service and control evidence. It should specify the limits and the review triggers. A stage label alone cannot carry that decision. The organization needs to know what it is relying on, what it would lose without it and who can change course when the original case no longer holds.

5. Redesign work without assuming expertise transfers automatically

AI can change the mix of tasks inside a role, but access to a tool does not establish competence in every task it makes possible. A person may produce a convincing draft outside their specialist knowledge without being able to judge its weaknesses. The transformation leader should identify where the new process requires review, training or a different allocation of expertise, rather than assume that technical capability automatically transfers to the employee.

Begin with the decisions people will make in the revised workflow. Some may move from production to review; others may handle exceptions, customer conversations or interpretation. These responsibilities should be explicit in the process design. If review is essential, the organization must allocate time and qualified people to it. Calling the workflow automated while relying on invisible specialist checking understates both the cost and the operating dependency.

Performance measures should follow the revised purpose. A person assigned to deeper customer work may complete fewer routine cases while creating more useful value. A reviewer may prevent an error that never appears in the production count. The program should decide how those contributions will be assessed, so that employees are not asked to change their work while being rewarded solely for the old activity measure.

Training should address the actual workflow and its failure modes. Generic tool instruction may be insufficient for deciding when an output is unreliable, when a task exceeds the approved scope or when to seek specialist help. Practice should include ambiguous and unsuccessful cases, not just demonstrations that show the tool at its best. The operating owner should know how people will obtain support after the initial rollout.

The organization also needs feedback that can challenge the design. If employees are penalized for reporting extra review or a task that does not work well, the transformation program loses important evidence. Managers should ask what changed in the real work and which assumptions failed. That information can support a narrower or better-designed process without requiring the team to abandon every useful capability.

The gate is a credible workforce plan for the changed process: responsibilities, support, evaluation and escalation. It is not a forecast that a particular share of jobs will disappear or a promise that every employee will move to higher-value work. Those outcomes depend on management choices and business conditions. The plan should make those choices explicit enough to review when the workflow expands.

6. Keep the transformation record current after launch

Production launch is an important milestone, but it does not end the transformation work. Users, demand, data and operating requirements can change. A workflow that met its purpose during a limited rollout may need a different design at larger scale. The transformation leader should establish which evidence remains current and which changes require a renewed decision.

The separate bottleneck instrument shows why a narrow deployment view is insufficient. Data access and quality is selected by 29 of 75 respondents, compute by 22, integration by 21, inference cost by 19 and governance and approval by 16. These are reported constraints in an operator sample, not a causal ranking for every enterprise. They point to several possible limits that can remain after a technical release. [BOTTLENECK]

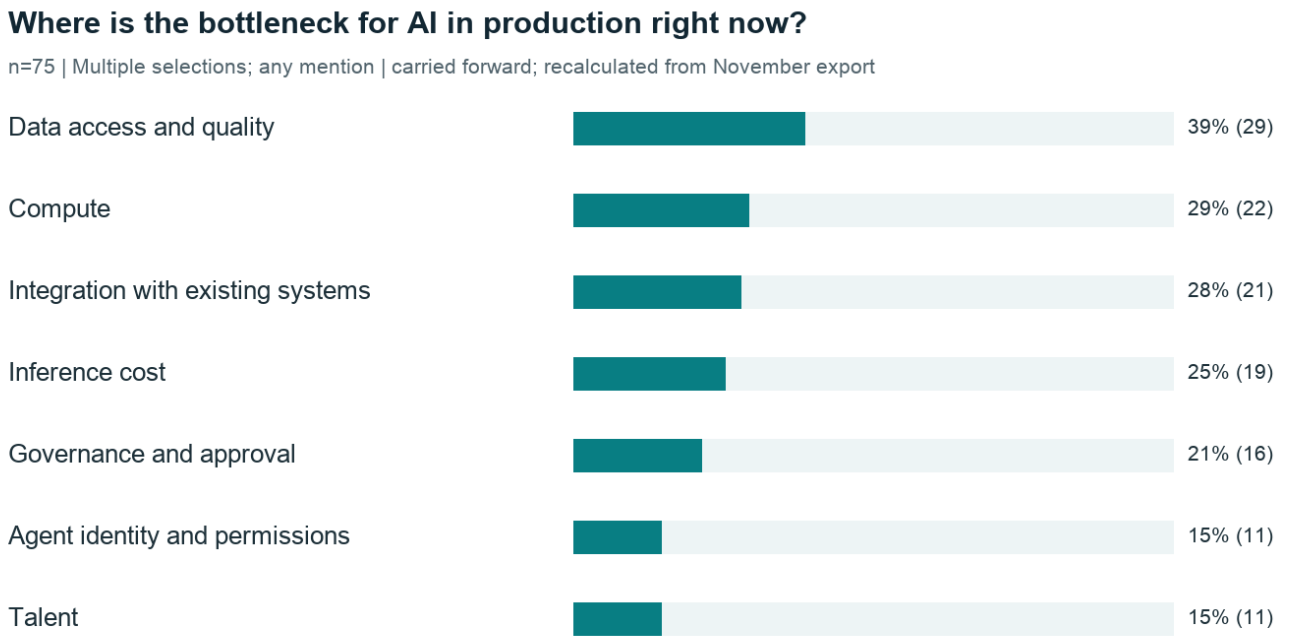
The review should follow the end-to-end outcome rather than just the application's activity. If the accepted service result deteriorates, investigate whether the cause lies in the model, source data, workflow, staffing or downstream process. An application metric can remain healthy while the business outcome worsens because another part of the system changed. The operating record should make those dependencies visible enough to test.

Evidence should be versioned. Preserve the scope, definitions and configuration that supported the original decision, then record material changes. A revised output definition or cost allocation can make a performance chart look different without an actual operating improvement. The program should distinguish measurement changes from business changes and avoid claiming a trend across incompatible definitions.

The review also needs an effective action path. If a benefit fails to appear, who can narrow the rollout or revise the process? If dependence becomes unacceptable, who can restore alternative capacity or renegotiate a commitment? A transformation dashboard that identifies issues without assigning decision rights merely documents delay. The CEO report provides the executive delegation framework; this report supplies the operating evidence that should trigger it.

The final gate is continuing ownership. The program can hand the workflow to normal operations when responsibilities, evidence and change rules are functioning. It should not keep a permanent transformation label simply because the technology is novel. The intended result is an operating process the business can manage, with a clear record of what changed and how its continued value will be judged.

Exhibit 2. Reported production bottlenecks



BOTTLENECK | Registration range: 2026-08-30 to 2026-09-08 UTC | Selective self-report sample

Exact selection	Count	Base	Any mention
Data access and quality	29	75	39%
Compute	22	75	29%
Integration with existing systems	21	75	28%
Inference cost	19	75	25%
Governance and approval	16	75	21%
Agent identity and permissions	11	75	15%
Talent	11	75	15%

Exhibit note: BOTTLENECK; carried-forward operator instrument, multi-select. These are stated constraints, not measured causes of a maturity stage or evidence of transition failure.

External evidence: the boundary between task help and organizational change

Stanford HAI's account of Karunakaran and colleagues' workplace research links engagement with both the purpose presented to employees and the practical support provided. [E-STANFORD] For transformation planning, the relevant implication is to examine whether training, exploration time and revised responsibilities support the promised change. The source does not isolate communication alone as a universal cause, and this report does not use it to predict a particular adoption rate.

Harvard Business School's institutional account of the GenAI Wall Effect describes how expertise and proximity to the task shape results. [E-HBS-WALL] That complicates a simple plan to redistribute work wherever a tool can generate an answer. The organization should identify who can evaluate the output and how the receiving process will use it. The report uses this qualitative boundary without importing unverified numerical estimates from the underlying paper.

METR's 2026 update explains why its newer developer study cannot provide a reliable current productivity estimate: participation and time measurement changed in consequential ways. [E-METR] A transformation program faces a related design challenge when volunteers, tasks or work patterns change between the pilot and rollout. The appropriate response is to inspect comparability and preserve uncertainty, not to choose the most convenient headline from an earlier or later study.

NIST's voluntary framework connects governance with understanding, measuring and managing AI risk. [E-NIST] In this report, that context supports treating the acceptance of dependence as an organizational decision. The proposed record is an OFF management tool, not a mandated template or a certification. The company still needs to apply its own policies, sector requirements and risk tolerance to the actual process.

The sources are useful because they address different missing steps. Employee engagement does not establish output quality. Task performance does not establish financial capture. A pilot comparison does not automatically remain valid after the work changes. A governance framework does not demonstrate that the specific control works. Transformation requires management to connect those steps with evidence rather than assume that one favorable result answers all of them.

There is also a limit to how much measurement can precede action. Some operating changes cannot be fully understood without trying them. The report therefore recommends bounded learning and progressive commitments, with stronger evidence required as exposure and dependence grow. The organization can act under uncertainty while remaining honest about what has been observed, what is inferred and what it has chosen to risk. That distinction is more useful than either demanding impossible certainty or declaring transformation complete after a successful demonstration.

Source	What it adds	What it does not establish
Stanford HAI workplace research account	Organizational support and engagement context	A universal adoption effect from framing alone

Source	What it adds	What it does not establish
HBS GenAI Wall Effect account	Expertise and task-boundary context	Reliable specialist performance in every new task
METR experiment update	Comparability and measurement cautions	A current universal developer-productivity estimate
NIST AI RMF	Organizational risk-management structure	Certification of an operating model

Applying the gates to a real operating decision

Use the review tool on one workflow before applying it to the entire portfolio. The transformation lead should convene the business owner, the receiving process owner and the functional specialists whose evidence matters. The first discussion should establish the intended outcome and the decision at hand. It should not begin by asking the team to defend a stage label already assigned in a dashboard.

Exhibit 3. Recommended evidence gates, separate from survey categories

Decision	Evidence required	Decision owner	Failure condition
Explore a process change	Problem, affected users, learning question and limit	Business sponsor	No meaningful decision would follow the result
Run a pilot	Defined output, comparison, support and next approver	Operating owner	Test cannot distinguish useful result from activity
Deploy	End-to-end acceptance, cost, controls and support	Operating owner with functional review	Material handoff or intervention remains unowned
Accept dependence	Benefit destination, reduced alternatives, continuity and change triggers	Accountable business executive	Commitment exceeds available evidence or authority
Transfer to normal operation	Current evidence, maintenance and escalation process	Service owner	No owner can sustain or revise the workflow

Consider an illustrative finance-close workflow. AI reduces the time needed to prepare an initial variance commentary, but the monthly management meeting still occurs on the same date and reviewers still wait for unresolved reconciliations. The preparation team reports a faster task. The business has not yet demonstrated an earlier decision or a shorter close. This scenario is invented to show the difference between a local improvement and a process outcome.

The first review maps the handoff. Which reconciliations delay acceptance? What evidence does the reviewer need? Can an earlier draft help identify issues sooner, or does it merely arrive before the supporting

data is ready? The next experiment might change the review sequence rather than the model. The team can preserve the drafting benefit while testing the actual constraint on completion.

Suppose the revised process produces an accepted commentary earlier and allows management to address a cash issue before the normal meeting. That is a more useful operating result, but the financial claim still needs care. The organization should document the decision and consequence rather than assign all subsequent cash improvement to AI. Other actions and conditions may matter. The CFO's benefit qualification process determines what can be accepted in the investment case.

Before reducing manual preparation capacity, the business should review dependence. Can the team complete the required work if the source or AI service is unavailable? What quality of fallback is acceptable, and for how long? Which knowledge must remain in the organization? The answer may justify a narrower fallback rather than a complete duplicate process. That choice should be made by the owner who is accountable for the finance service.

Working tool: transformation review record

Field	Required content	Evidence owner
Intended process change	Before/after workflow and customer or internal decision	Business owner
Accepted outcome	Definition, quality and downstream usability	Receiving process owner
Comparison	Period, task mix, support and relevant changes	Transformation / analytics lead
Benefit destination	Service, growth, capacity or financial conversion	Operating owner and finance
Workforce change	Responsibilities, training, review and evaluation	Functional leader
Dependence accepted	Alternatives reduced, fallback, obligations and limits	Accountable executive
Change trigger	Event that reopens the decision; next review	Service owner

The review should record one of several actions: expand, continue bounded learning, narrow the scope, redesign the handoff, defer pending a dependency or retire. Each action needs an owner and a condition for reconsideration. A decision to defer is useful only if it identifies what must become true; otherwise, it leaves the project consuming attention without a credible path forward.

At the next review, compare the actual result with the original learning or operating commitment. Keep the original definitions and note changes. If the workflow is useful but a larger benefit has not appeared, preserve the useful result and revise the broader claim. The transformation program should be able to tell

that mixed story without treating every qualification as failure. That is how the evidence improves the operating model rather than merely supporting a launch announcement.

The next transformation review

Choose one process where the organization is considering deeper dependence. Ask the business owner to explain what would change in cost structure, staffing plans or service obligations if the workflow became unavailable. This question follows the survey's embedded definition while turning it into a concrete internal decision, rather than assuming that the label itself establishes readiness.

Then inspect the process around the AI task. Identify the receiving team, the next decision and the constraint that could prevent the local improvement from becoming a useful result. Assign the complementary change to an owner with the authority and resources to make it. A transformation program cannot capture a benefit that remains nobody's operating responsibility.

Finally, decide which commitment the current evidence supports. It may justify a production release, a narrower scope, more learning or a deliberate acceptance of dependence. Record the limits and the conditions that would reopen the decision. Do not require every workflow to reach the same stage or treat retirement of an unsuitable experiment as a failure to transform.

The November baseline remains a cross-section of 91 answers. The contribution of this report is the decision process built around that limitation: identify the work that changes, establish the result, assign the benefit and accept dependence with evidence. That gives the organization a way to manage transformation after the novelty of the tool has passed.

Reader questions

What does embedded mean in the survey?

The exact option is Embedded (removing it would change our cost structure or hiring plan). Twelve of 91 respondents selected it. [MATURITY-04]

Did more functions become embedded in October?

The supplied data does not establish that. The direct maturity distribution is unchanged and no comparable new answers were added.

Is 54/91 a transformation success rate?

No. It combines deployed and embedded answers in one cross-section. It does not measure completion of a journey, value realization or control effectiveness. [MATURITY_DEPLOYED_OR_EMBEDDED]

Are the evidence gates survey stages?

No. They are a recommended management process, kept separate from the four response categories and from market-map advisory models.

Must every production workflow become embedded?

No. A bounded production use can be the appropriate destination. Deeper dependence should be justified by the business purpose and evidence.

What if a pilot improves one task but not the whole process?

Preserve the local finding and inspect the surrounding constraint. Do not claim an end-to-end benefit until the relevant handoffs and outcomes support it.

Does a fallback plan prove that dependence is acceptable?

No. Management must assess whether the fallback's actual service level and limitations fit the process. Documentation alone does not establish adequacy.

Citable facts

- Exploring 19, piloting 18, deployed 42 and embedded 12, base 91, carried-forward single-selection maturity instrument. [MATURITY]

- Deployed or embedded: 54/91, 59%, a cross-sectional union, not a transition rate. [MATURITY_DEPLOYED_OR_EMBEDDED]

- Data access and quality: 29/75, 39%, reported production bottleneck in a separate operator question. [BOTTLENECK-01]

Methods and limitations

Calculation-v1.1 selects eligible non-invited, nonblank answers within the cutoff after excluding byte-identical exports. Deduplication is separate per question by normalized email, ordered by registration timestamp, filename and row. Registration time is not an answer-edit timestamp. The maturity instrument asks about the respondent's function, not independently verified enterprise maturity.

MATURITY is single selection, n=91, from an export with 314 eligible application rows. BOTTLENECK is multi-select, n=75. MKT is separate updated cumulative context, n=314; no matched relationship to maturity is asserted. The report does not infer transitions, stage duration, causation, market representativeness or a composite transformation index.

The evidence gates and dependence record are management recommendations. The finance-close scenario is illustrative. The HBS institutional account is used only for qualitative context; exact numerical estimates from the unavailable underlying paper are not repeated. No private respondent quotations are included.

References

[OFF-DATA] Open Future Forum. Executive AI Statistics, November 2026; aggregate data and source manifest.

[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-HBS-WALL] Harvard Business School AI Institute. Why AI Helps Until It Doesn't: Inside the GenAI Wall Effect. 18 September 2025. https://aiinstitute.hbs.edu/why-ai-helps-until-it-doesnt-inside-the-genai-wall-effect/

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

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

[OFF-CONTINUITY] Open Future Forum. 30-Day AI Continuity Drill, October 2026. https://openfutureforum.com/blog/30-day-ai-continuity-drill

[OFF-PRODUCTIVITY] 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 CFO report qualifies financial benefits; AI Leaders owns production release evidence; CEO owns allocation and delegation; Board Preview covers material dependence and oversight.

Suggested citation: Newlands, Murray. AI Transformation Report, November 2026, Edition 4. Open Future Forum. Evidence cutoff 31 October 2026.

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

Editable working tool

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