NOVEMBER 2026 | EDITION NO. 03

THE EXECUTIVE FRUSTRATION BRIEF

What business leaders are struggling with - and what it means.

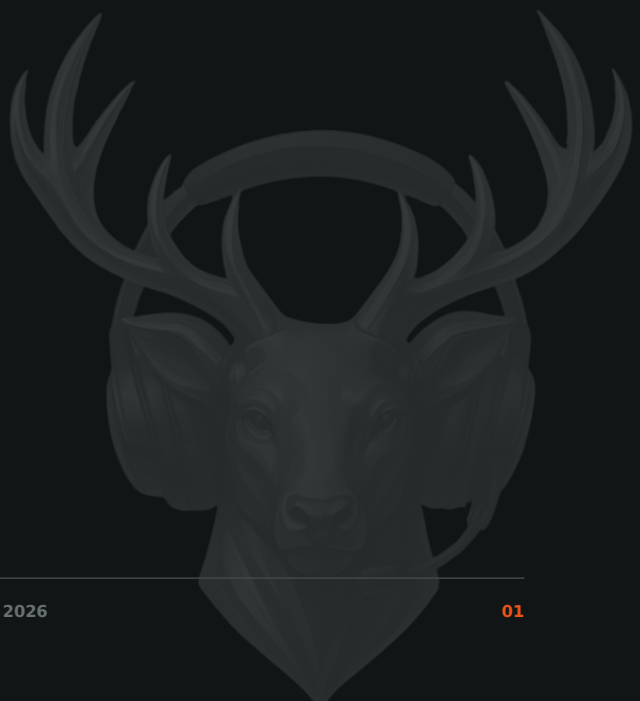
AI is moving faster than the operating proof around it.

Value, authority, cost and human work are becoming one executive decision.

13%

of CIOs report significant or transformational business value from AI. Gartner survey of 2,002 leaders [1]

Reported value, not independently audited financial return.



OCTOBER EVIDENCE. NOVEMBER DECISIONS.

The operating proof gap.

September named the execution debt: AI capability added faster than work was redesigned. October showed the operational consequences. November asks for proof that the changed workflow is worth its cost, authority and human effort.

Across five October research cycles, financial proof, agent authority, workforce redesign and consumption economics formed the leading priority tier. These are connected decisions. One operating record should link the expected result, the full workflow cost, the permissions granted, the exceptions people carry and the decision to scale, redesign or stop.

AI working is not enough. Can leadership prove the changed workflow earns its full cost?

What genuinely sharpened in October

COST PER ACCEPTED RESULT

Consumption economics recorded the largest editorial increase (+20 points). Runtime, retries, review and recovery belong in funding decisions.

AUTHORITY YOU CAN REVOKE

New NIST guidance makes identity, delegated scope, live revocation and retained evidence part of the operating design. [3]

WORK AFTER AUTOMATION

Harder exceptions, supervision and correction reduce the net capacity actually released by AI.

CUSTOMER RECOVERY

Human access is only useful when context transfers and the next person has authority to resolve the issue. [7]

What held: Financial proof and agent accountability stayed prominent. Risk and SMB payback received less fresh evidence late in the month, but that does not mean either problem was resolved.

October movement is comparative ParkWest editorial priority, not a statistical change in market prevalence. The opening baseline is Sept. 25-Oct. 1; the closing week is Oct. 23-29.

The three-edition trend monitor.

The storyline has progressed from the underlying problem, to its operational consequences, to the executive decision. Each cell summarizes editorial emphasis in that edition, not a measured change in market prevalence.

THEME	SEPTEMBER	OCTOBER	NOVEMBER
Value & financial proof	Spend grows; blind spots	Baseline a workflow	Stop rules and unit cost
Agent authority	Autonomy beyond policy	Permissions in practice	Identity and revocation
Customer trust	Automation can fail	Human handoff leads	Recovery quality
Workforce capacity	Supervision emerges	Net capacity matters	Roles and exceptions
Consumption cost	An emerging exposure	Less fresh evidence	Largest index mover
SMB implementation	Selective payback	Bounded use cases	Core integration gap

OCTOBER: THE FOUR-WEEK ARC

OCT 2-8
Accountability

OCT 9-15
Proof before scale

OCT 16-22
Stop criteria

OCT 23-29
Cost and recovery

The progression: Execution debt was the diagnosis. Handoffs exposed its consequences. October's newest evidence now makes it a decision about accepted outcomes, enforceable authority and stopping rules.

Trend cells compare the September, October and November ParkWest editions; different underlying surveys cannot be plotted as one statistical series. The separate October index uses a consistent editorial baseline.

Priority is converging. Momentum still differs.

October scores cluster near the maximum. For the public view, bands distinguish closing priority while index-point movement shows what changed. The underlying scores remain in the research record.

LEADING 100 **ELEVATED 95-99** **WATCH BELOW 95**

EXECUTIVE FRUSTRATION	BAND	MOVE	READING
AI spend without defensible value	**Leading**	**+3**	Persistent
Agent autonomy and permissions	**Leading**	**+2**	Persistent
Workforce redesign and capability	**Leading**	**+3**	Exception load
AI consumption cost at scale	**Leading**	**+20**	Largest mover
Customer trust at failed handoffs	**Elevated**	**+1**	Late rise
Data and integration debt	**Elevated**	**+9**	Stronger
Audit and assurance	**Elevated**	**+10**	Traceability
SMB adoption and fuzzy payback	**Watch**	**+3**	Less late evidence
Third-party AI / identity risk	**Watch**	**-2**	Priority shifted
Sales pipeline reliability	**Watch**	**+13**	Cooled late

READ THE MOVEMENT, NOT JUST THE RANK

+20 for consumption economics means stronger ParkWest editorial urgency since the opening baseline - not a 20% increase in AI prices. A category can rise over the month while its final week cools because fewer fresh signals appear.

Score changes compare Sept. 25-Oct. 1 with Oct. 23-29. Six editorial criteria produce a 0-100 index. Bands are derived from closing scores; neither scores nor bands measure market prevalence.

Measure the accepted result, not the AI bill.

Gartner found that only 13% of surveyed CIOs reported significant or transformational AI value. Only 31% had criteria to stop a pilot, and 17% had criteria to end low-value production systems. [1]

71% **track total AI inference and API spend**

44% **track per-unit inference cost. [2]**

The missing measure is often full cost per accepted business outcome - including retries, review and exception recovery, not just the model or software bill.

ILLUSTRATIVE WORKFLOW / NOT A CASE STUDY

100 attempts **80 accepted outcomes**

$35	$20	$90	$55
AI runtime	Retries	Human review	Recovery

$200 total / 80 accepted = **$2.50**

Test one funded workflow: agree on the baseline, quality threshold, fixed setup cost, runtime/retry cost, human review, recovery, accountable owner and decision date. Compare against today's end-to-end workflow at the same quality standard.

THE EXECUTIVE QUESTION

What does one accepted result really cost - and what result would make us stop?

Example figures are hypothetical and exclude one-time implementation cost. Time released is not cash saved unless an expense actually changes. Source surveys are self-reported; stopping-rule associations are not causal proof.

An autonomous action needs an accountable trail.

Deloitte found that 40% of surveyed organizations allow AI to act without human approval to at least a moderate degree. For 23%, accountability after a negative outcome remains unclear. [2]

NIST's October design tenets turn the question into a workflow test: distinct identity, delegated scope, time-limited authorization, human-review triggers, monitoring, evidence and revocation. [3]

Five questions that turn policy into evidence

01	INTENT	What human request or approved business result initiated this action?
02	IDENTITY	Which specific agent acted - and on whose behalf?
03	SCOPE	What could it read, change, send or spend, and for how long?
04	EVIDENCE	Can we reconstruct tool calls, approvals and resulting changes?
05	RECOVERY	Can a named owner revoke authority and contain or reverse harm?

WHAT NOVEMBER ADDS

September raised governance. October examined enforceable boundaries. November asks for a defensible record of each consequential action. BCG also finds that higher-accountability organizations are more likely to deploy kill switches, an association rather than causal proof. [8]

THE EXECUTIVE QUESTION

If an agent changes a record incorrectly, who can trace it, stop it and repair it?

NIST tenets are proposed design guidance, not evidence of deployment. BCG results are associations from a global enterprise sample; they do not demonstrate that kill switches prevent incidents.

The work left behind is becoming visible.

The strongest workforce signal was not simply that AI may remove tasks. It was that automated work often leaves harder exceptions, more review, broader roles and new coordination behind.

4% **of surveyed CIOs reported explicit plans to redesign roles for a human-agent workforce. [1]**

Charter's study of 5,002 U.S. desk workers, managers and executives found that employee agency - influence over how AI fits into the job - was the strongest enabling factor for reported benefit. Executives were also more likely than workers to trust AI output. [5]

Ricoh separately found a gap between business decision-makers' confidence in workplace technology and employees' experience, reinforcing the risk of designing from the executive view alone. [6]

Gross time saved is not removable cost.

Map one role into repeatable tasks, exceptions, decisions requiring judgment, review and correction, learning time and coordination. Give recovered capacity an explicit destination: backlog, customer follow-up, quality, growth or risk work.

The executive question

What work will people do differently next month - and have we given them the time, authority and practice to do it well?

The cited workforce results are self-reported associations. They do not establish that AI use caused the reported outcomes or that time savings equal headcount reduction.

Customers are not rejecting AI. They are rejecting dead ends.

October's customer evidence was consistent: people accept automation when it resolves the issue, but trust falls when the path hides a person, drops context or produces an answer no one can defend.

92% **of consumers in Five9's pooled U.S., U.K. and German sample gave AI less than 10 minutes before asking for a person. [7]**

The same research reported that trust more than doubled when a human option was visible. The number should not be generalized to every market or interaction, but the mechanism is commercially relevant: an escape path is part of the product experience. [7]

The opportunity is not maximum containment. It is dependable resolution across the automated path and the failed path.

An opportunity worth testing

Choose one journey with repeat contacts, complaints or policy exceptions. Test the normal path and the failed path. Record whether context transfers, whether the person has authority to resolve the issue, and whether the final answer is supported by trusted policy evidence.

Judge the whole journey on completed resolution, repeat contact, customer effort, abandonment, recovery time and total cost. A lower bot-containment rate may be acceptable if the customer and the economics improve.

The executive question

When the automated path cannot finish the job, does the customer reach someone who understands the problem and can act?

Five9 sponsored the study; fieldwork occurred in April 2026. Reported preferences do not establish causal retention or revenue lift.

Experimentation is not implementation.

Goldman Sachs reported that 86% of surveyed 10,000 Small Businesses alumni use AI, but only 15% have fully integrated it into core operations. The sample is program alumni, not a representative estimate for all U.S. small businesses. [4]

15% **reported full integration into core operations.**
Survey of 2,473 program alumni. [4]

This is not proof that every small business needs an AI transformation. It is evidence of a practical gap between trying tools and changing how the business operates. A useful starting point is smaller: one recurring loss, one decision date and one owner who can act on the result.

Start with the problem, not the category.

Compare at least four options: leave the workflow alone, simplify the process, use an existing tool or automation, or add AI with the appropriate human control. The correct answer may be no purchase.

A 90-day proof window can structure the decision, but it is not a promised payback period. The test should state the baseline, full cost, accepted outcome, review date and the condition for stopping.

The executive question

What is the one recurring operating problem worth solving before we add another tool?

The Goldman Sachs sample is useful directional evidence from program alumni and should not be generalized as national SMB prevalence.

Bring me one business problem - and the decision date.

I do not need a list of tools. Show me where the work breaks, who carries the exception, what the customer experiences, what the result costs and what happens when the system is wrong.

If we can answer those questions, we can decide whether AI belongs in the workflow. If we cannot, another demonstration will not make the decision clearer. The right first move may be a process change, better data, a clearer handoff, a different owner - or no technology purchase at all.

Show me the accepted outcome. Show me who can act. Show me the work people still do. Show me what makes us stop.

That is where October's evidence leads me. The business advantage is not the largest AI portfolio. It is the ability to make a clearer decision, prove what changed and redirect quickly when the workflow does not earn its cost.

SIMMELLE ZANE | PARKWEST-CX

Six situations worth bringing.

01 **We use AI but cannot defend the spend.**

Reconcile one workflow to its full cost, accepted outcome, owner and decision date.

02 **The agent can act but the limits are unclear.**

Define identity, delegated scope, review threshold, retained evidence and shutdown authority.

03 **AI saves time but the workload does not fall.**

Measure review, exceptions, coordination and the destination of recovered capacity.

04 **The bill grows as usage grows.**

Model cost per accepted outcome and authorize a response threshold before scale.

05 **The customer repeats the issue or disputes the answer.**

Repair the knowledge path, context transfer, human authority and recovery measure.

06 **We have many pilots and no confident decision.**

Set criteria to scale, redesign, defer or stop before the next funding gate.

One problem. A clearer decision.

These are decision scenarios derived from the monthly evidence, not promised outcomes or case studies.

What we will watch in November.

01 **Stop rules that are actually used**

Whether organizations redirect low-value AI work when a funding or production decision arrives.

02 **Cost per accepted outcome**

Whether finance teams connect runtime usage, retries, review and recovery to workflow economics.

03 **Agent revocation in practice**

Whether named owners can interrupt authority and retrieve evidence under real conditions.

04 **Capacity that appears in the business**

Whether saved time becomes backlog reduction, growth, quality or lower cash cost.

05 **Resolution after the failed path**

Whether visible human access and retained context improve final customer resolution.

The question I would ask your leadership team

Which AI-supported workflow would we stop funding today if we had to defend its full cost, authority, human work and accepted outcome?

YOUR STARTING POINT

Bring ParkWest one business problem. The complimentary 10-minute executive decision experience produces an Executive Action Brief and helps clarify the next move before the tool decision.

parkwest-cx.com/assessment

Evidence, comparisons and limits.

Five October weekly research cycles underpin within-month movement. September and October ParkWest editions supply qualitative historical context. The source titles below link to the original publishers.

[1] GARTNER | OCT. 19, 2026

2027 CIO and Technology Executive Agenda

2,002 technology executives; fielded May-July.

[2] DELOITTE | OCT. 29, 2026

State of AI in the Enterprise

3,500 senior leaders globally; May-June.

[3] NIST | OCT. 27, 2026

Design Tenets for Agent Identity Management Systems

Proposed agent identity and authorization design.

[4] GOLDMAN SACHS | OCT. 22, 2026

The AI Economy Is a Small Business Economy

2,473 alumni of 10,000 Small Businesses; August.

[5] CHARTER | OCT. 21, 2026

Fall 2026 AI Pulse

5,002 US AI-using workers; plus 91 interviews.

[6] RICOH USA | OCT. 21, 2026

2026 Workplace Technology Adoption Gap

2,017 employees and decision-makers; US/Canada.

[7] FIVE9 + HANOVER | OCT. 14, 2026

2026 Consumer CX Report

3,000 consumers; US, UK and Germany; April.

[8] BCG | OCT. 26, 2026

How to Create AI Value by Balancing Innovation and Risk

1,330 leaders of global enterprises.

METHOD / HOW TO READ THIS SERIES

> **Facts:** external study findings are attributed; publication month does not equal fieldwork month.

> **Interpretation:** the three-edition trend monitor describes changes in editorial focus - not a longitudinal statistical sample.

> **Index:** ParkWest scores six 1-10 editorial factors (evidence, impact, AI relevance, ParkWest fit, commercial trigger and momentum); month movement compares the opening and closing weekly baselines.

> **Decisions:** recommended tests and the cost example are proposed operating designs, not validated customer results or promised payback.

Comparable source populations, denominators and measures are required before external survey results can be presented as a statistical month-over-month trend.