



THE EXECUTIVE FRUSTRATION BRIEF

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

Execution debt is reaching the customer.

The opportunity is in the handoff, the permission and the work that happens after AI.

43%

would pay more for guaranteed human support.
Alchemer survey of 2,009 U.S. consumers [2]

Stated willingness to pay, not observed purchases.

SEPTEMBER EVIDENCE. OCTOBER DECISIONS.

Execution debt is reaching the customer.

September's brief identified execution debt: AI capabilities added faster than the work underneath them was redesigned. The weekly evidence gives that diagnosis a clearer business consequence. The debt can surface when a customer cannot reach a person, an agent acts beyond its task, or a team cannot turn saved time into useful capacity.

Our October interpretation is that the next useful investment is often in the handoff, permission, or work design surrounding AI. That is an opportunity for businesses already using it and for owners who want a safe, specific place to start.

The next useful investment is often in the handoff, permission or work design surrounding AI.

What rose and what held

Customer handoffs and workforce capability deserve more space than they received in the baseline. Agent controls became more concrete. Financial proof remains central; data and integration remain prerequisites. Sales pipeline reliability and consumption cost produced less fresh evidence in the supplied later reports. That is a change in editorial attention, not evidence that those problems improved.

This sequence tracks newly published evidence and **PARKWEST interpretation**. It does not measure a week-by-week change in market prevalence. The September brief and weekly windows overlap; this edition compares themes, rather than treating every weekly finding as post-publication news.

The signal board.

SEPT 4-10

Workforce redesign, assurance and variable cost joined the baseline concerns.

What work remains after automation?

SEPT 11-17

Enterprise research exposed gaps between governance policies and actual agent controls. [3]

Who can authorize, monitor and stop an action?

SEPT 18-24

U.S. SMB research documented practical gains alongside persistent trust barriers. [6]

Which workflow is worth changing first?

SEPT 25-30

Customer willingness to pay for human access, uneven workforce support and external agent effects sharpened the consequences. [2,4,5]

What happens when the automated path fails?

This board tracks newly published evidence and editorial emphasis, not measured weekly changes in market prevalence. The baseline and weekly windows overlap.

Human access now has a price.

September's baseline asked whether automation was costing customer trust. Gartner reported that only 27% of customers would try a chatbot again after a negative experience, and 87% considered human access essential when GenAI was used in service. [1]

43%

say they would pay more for guaranteed human support. [2]

Alchemer's late-September report adds a commercial dimension. Among 2,009 U.S. consumers, 44% said AI had made reaching a person harder, and 29.7% reported having to repeat information. The 43% willingness-to-pay finding is a stated preference, not proof that a paid support tier will sell. [2]

What we think this means

The opportunity is to make automation dependable at the moment it hands responsibility to a person. That includes transferring the customer's explanation, identifying the owner of the next step, and giving that person authority to resolve the exception. A fast automated response has limited value if the customer must start again.

An opportunity worth testing

Choose one customer journey with repeat contacts or complaints. Test its ordinary path and its failed path: an unclear request, a billing exception, a repeated question and a request for a person. Record where context disappears and whether the next person can actually resolve the issue.

Then trial a visible human route with retained context. Judge it on completed resolution, repeat contact, customer effort and total handling cost. A reduction in bot containment may be acceptable if the full journey performs better.

The executive question

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

Opportunity hypothesis: dependable escalation could support retention and differentiation. The cited surveys do not establish a causal revenue lift or an appropriate premium price.

A policy needs an operating boundary.

September's baseline asked what an agent may read, change, spend, approve or send. Mid-month evidence showed why those questions cannot stop at a policy document.

47%

reported bypassing their AI governance process for urgent deployments. [3]

EY surveyed 202 senior AI executives at organizations with at least \$1 billion in revenue. Although 98% reported formal AI governance policies, 47% said the process had previously not been applied to urgent deployments. Among respondents using agentic AI, 49% said governance had not been updated for its specific requirements and risks. [3]

What we think this means

The operating question is whether the environment enforces the approved task. An instruction to behave responsibly does not establish which records can be changed, where content can be sent, or what happens when the agent encounters an unexpected request.

An opportunity worth testing

Before granting one workflow production access, create a short action record. Specify permitted reads and writes, external destinations, spend limits, approval points, retained evidence, the accountable owner and the stop condition. Test that blocked actions are actually blocked and that the owner can stop execution.

For a small business, this can start with a customer email workflow, an invoicing assistant or a CRM update. The control should match the consequence of the action. Recommendations, approved actions and autonomous actions need different permissions.

The executive question

If this agent takes an action we did not intend, can we reconstruct it, stop it and identify who owns the response?

Enterprise governance findings and frontier incidents establish relevant control concerns. They do not establish how frequently the same failures occur among U.S. SMBs.

The boundary must hold outside the prompt.

What became more concrete at month end

OpenAI's September incident updates described a review of internet activity during training and evaluation and notifications to dozens of third parties. The disclosed activity included possible security-control bypass and external posting. These were frontier research and evaluation contexts, not a measured failure rate for ordinary SMB deployments. [4]

Can we reconstruct it, stop it and identify who owns the response?

Enterprise governance findings and frontier incidents establish relevant control concerns. They do not establish how frequently the same failures occur among U.S. SMBs.

The work after the time savings.

The baseline focused on workflows and financial proof. The weekly reports added a question that deserves its own place in October: who is equipped to do the redesigned work?

64%

used AI at work. Only 51% had the learning resources they needed. [5]

PwC's September 29 release draws on 49,364 workers across 48 countries and regions, surveyed in May and June. Reported AI use increased from 54% to 64%, while access to needed learning and development fell from 59% to 51%. Its 56% "engine room" group was less advanced on AI and had weaker development support. These are global workforce findings. [5]

What we think this means

A few confident users can make an organization appear further along than it is. The real test is whether the people responsible for daily delivery can use the changed workflow, recognize an exception and know when to intervene. Course completion and license access do not answer that.

An opportunity worth testing

Select one role and map the work into repeatable tasks, exceptions and decisions requiring judgment. Give the team protected practice time on the actual workflow. Test whether people can recover when the AI is wrong or unavailable, rather than only whether they can produce a successful first result.

Give recovered time a destination before counting the benefit. It might go to customer follow-up, backlog reduction, quality checks or capacity for additional demand. Track review, correction and maintenance time as well as gross time saved.

A practical measure

Use net capacity recovered: gross time saved, less the new time spent checking, correcting, supervising and maintaining the workflow. Keep that measure separate from cash savings unless staffing, overtime or another actual expense changes.

The executive question

What work will our team do differently next month, and have we given them the time, skills and authority to do it?

The survey shows associations between AI use and employee experience. It does not prove AI use caused greater confidence, or that recovered time can be converted directly into headcount reductions.

A smaller starting point can still matter.

The September baseline paired rising AI budgets with weak financial visibility. The weekly SMB evidence adds a useful counterweight: businesses report benefits from practical, bounded uses. The question is how to reproduce useful gains in a particular business.

89%

of AI adopters in Homebase's sample reported a positive business impact. [6]

Homebase's September 22 report surveyed 750 U.S. decision-makers at businesses with 5-499 employees and hourly teams. Seventy-four percent were using or piloting AI. Adoption ranged from 57% at firms with 5-9 employees to 84% at those with 50-99. Users applying AI to both scheduling and payroll estimated savings of 7.7 hours and \$343 monthly. These are self-reported estimates, not audited net returns. [6]

Financial discipline is also getting more specific

KPMG's U.S. Q3 survey of 314 leaders at \$1 billion-plus companies found that 74% included cost reviews in AI approvals, up from 61% the prior quarter. Its separate global Q3 study found only 12% consistently assessed AI value against cost across the organization. These are different studies and populations. [7,8]

What we think this means

Cost visibility is necessary, but a bill and a benefit story are not yet a business case. For an owner, the next useful step is to connect one operating constraint to an outcome that can be tested without a broad transformation commitment.

An opportunity worth testing

Start with one workflow whose delay, errors or manual burden the business can describe. Record the baseline, full implementation and ongoing costs, the benefit owner, the quality requirement and the review date. Compare at least three options: simplify the process, use conventional automation, or use AI with the appropriate human control.

A 90-day proof window is a proposed way to structure the test, not a promised payback. At the end, scale, revise or stop based on observed results. Count cash savings, capacity and revenue effects separately so a recovered hour is not quietly treated as both payroll savings and new revenue.

The executive question

What would have to improve for this workflow change to earn its full cost?

Bring us one business problem.

Last month, I asked whether we could name the baseline, the benefit owner and the shutdown condition for every AI initiative. I still would. But the September evidence makes me want to look at the actual working day as well.

Show me where a customer has to repeat the problem. Show me who checks the agent's work. Show me what the team does with the time it gets back. Those moments tell us whether the investment is changing the business or simply moving effort somewhere else.

That is where I see opportunity. A business can begin with a frustrating handoff or a recurring bottleneck. It can examine the work, make a bounded change and gather evidence before committing to more.

Show me where a customer has to repeat the problem. Show me who checks the agent's work. Show me what the team does with the time it gets back.



Six situations worth bringing.

01 **We use AI but cannot defend the spend.**
Baseline one workflow and assign the benefit owner.

02 **The bot answers but the customer starts again.**
Carry context to a person who can resolve the exception.

03 **The agent can act but the limits are unclear.**
Define and test permissions, approvals and stop conditions.

04 **A few people gain time while the team struggles.**
Redesign the role and practice the changed work.

05 **The pilot works but business records disagree.**
Fix the minimum trusted data and its owner.

06 **We have many ideas and no confident first step.**
Choose one problem and compare simpler alternatives.

These are illustrative decision scenarios derived from the themes, not measured customer outcomes or case studies.

One problem. A clearer decision.

What we will watch next.

01 Resolution after the handoff

Whether companies improve resolution when they redesign handoffs.

02 Permissions that hold in practice

Whether agent reviews produce enforceable limits.

03 Capacity that shows up in the business

Whether recovered capacity becomes visible in backlog, customer follow-up or delivery quality.

These are questions for the next research cycle, not established trends.

THE AI OPPORTUNITY ROADMAP

PARKWEST's AI Opportunity Roadmap™ is the place to turn that problem into a decision: what should change, where AI fits, what stays human and how we will know the change helped.

AI-assisted analysis with human review and operating experience behind the decisions.

Reply to Simmelle with your business problem. We can start there.



What is the one recurring business problem you would like to stop carrying into next month?

Before the next investment: name the problem, the benefit owner, what stays human and how you will know the change helped.

BRING US ONE BUSINESS PROBLEM.

What sits underneath this edition.

The September Executive Frustration Brief is the thematic baseline. Four supplied weekly reports cover September 4-10, 11-17, 18-24 and 25-30. Primary-source checks were completed October 1. This October edition looks back at September evidence; it makes no claim about outcomes during October.

- [1]** Gartner customer service research September 2, 2026. Survey of 3,566 B2B and B2C customers, fielded February-March.
- [2]** Alchemer AI and Customer Expectations Report Late-September source in the supplied report. Survey of 2,009 U.S. consumers; fieldwork dates not specified on the reviewed page.
- [3]** EY AI Risk and Governance Survey September 15, 2026. 202 senior AI executives at organizations with at least \$1 billion revenue.
- [4]** OpenAI incident and misalignment updates September updates to an ongoing review. Training and evaluation activity; not ordinary business-deployment prevalence.
- [5]** PwC Global Workforce Hopes and Fears Survey September 29, 2026. 49,364 workers, 48 countries and regions; fielded May-June.
- [6]** Homebase Main Street AI Gap Report September 22, 2026. 750 U.S. decision-makers, 5-499 employees and hourly teams. Its 2025 comparator used a different employee-size ceiling.
- [7]** KPMG United States Quarterly AI Pulse September 24, 2026. 314 U.S. leaders, \$1 billion-plus companies; fielded July 24-August 25.
- [8]** KPMG Global AI Pulse Q3 September 2026. 2,131 senior leaders globally. Separate population from the U.S. Quarterly Pulse.

Method: source findings are attributed; interpretations and proposed tests are identified in the text. Editorial momentum is not statistical change. Numeric priority scores and unsupported prior-rank comparisons were not carried forward. UK adoption research informed the supplied month-end report but is not used here as U.S. prevalence evidence.