



THE EXECUTIVE FRUSTRATION BRIEF

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

85%

PLAN TO INCREASE
AI SPENDING

vs.

22%

HAVE SUCCESSFULLY
SCALED AI

So what's actually working?

INSIDE THIS EDITION
THE EXECUTION DEBT
NOBODY BUDGETED FOR

Budgets are climbing. Visibility isn't.

Gartner surveyed functional leaders about AI. Eighty-five percent plan to increase spending in 2026. About eleven percent could not say what their function spent on AI in 2025.

11%

of functional leaders **could not say what their own function spent on AI last year** — while 85% plan to spend more this year.

SOURCE · GARTNER, SEPTEMBER 1, 2026

Sit with that pairing for a second. It is not a technology finding. It is a financial control finding. In most companies, an eleven percent blind spot on any other line item would trigger a conversation with the audit committee.

**“It is not a technology finding.
It is a financial control finding.”**

The same research found only 22% of organizations had actually scaled AI across multiple business units. So the picture is: budgets climbing, visibility patchy, and roughly one in five companies having genuinely operationalized any of it.

The interesting part is what separated the winners. Gartner's high performers

weren't distinguished by better models or bigger budgets. They measured returns continuously, moved money between initiatives, and killed the ones that weren't working. That's not an AI capability. That's portfolio management applied to something most companies have been treating as an experiment.



Companies Bought AI. They Skipped the Operating Model.

Look across the latest evidence and the frustrations that appear unrelated turn out to be the same frustration wearing different clothes.

The CFO who can't tie AI spend to a financial outcome. The CIO watching agents get access to production systems faster than anyone can define what those agents are permitted to do. The COO whose chatbot is deflecting contacts and losing customers. The CRO staring at a pipeline full of activity and no confidence in the forecast. The RevOps lead whose CRM technically contains everything and answers nothing.

These are not five problems. They are five symptoms of one thing: companies added capability faster than they redesigned the work underneath it.

EXECUTION DEBT

I'd call it execution debt. Pilots, tools, automations, dashboards, integrations, and now agents — layered on top of workflows nobody has intentionally redesigned in years. Each layer works in isolation. The system doesn't.

Watch where the money is going and this becomes hard to argue with. Google Cloud and Accenture are placing up to 1,000 forward-deployed engineers directly with customers. That is not a licensing move. Enterprises are paying for people to sit inside their business and make things actually work. Meanwhile vendors like Boomi and SAP are pitching agent control planes and integration governance — permissions, identity, token spend, auditability, connections to core systems. Those are not model features. Those are operating-model features.



The market has quietly stopped selling intelligence and started selling plumbing, permissions, and proof.

Here's the part executives find uncomfortable. If your AI ROI is unclear, the diagnosis is rarely "we picked the wrong model." It's usually that nobody baselined the workflow before the pilot started, nobody owned the benefit, and the definition of success became adoption. Adoption is not an outcome. It's an input that costs money.

1,000

forward-deployed engineers being placed **directly inside customer organizations** by Google Cloud and Accenture — capacity to execute, not capacity to license.

SIGNAL · ENTERPRISE IMPLEMENTATION CAPACITY



The five frustrations that mattered most.

The highest-pressure executive problems, ranked by how hard they are pushing on decisions right now.

01

Spending on AI with no defensible financial story

ACCELERATING

Executives are approving more AI investment while lacking a credible way to say which use cases improved economics.

WHY IT MATTERS — Capital stays trapped in pilots and licenses outlive the experiments that justified them.

02

Agents gaining autonomy faster than anyone can govern them

ACCELERATING

Companies can deploy agents that take action but haven't defined identity, permissions, spend limits, or who owns the consequences.

WHY IT MATTERS — The risk moved past hallucination into unauthorized action, data exposure, and audit failure.

03

Data and integration debt killing projects at production

RISING

The demo works. The rollout stalls because the customer, financial, and operating data is fragmented and context-poor.

WHY IT MATTERS — You cannot let an agent act on a record the business doesn't trust.

04

Busy sales teams, unreliable pipeline

RISING

Leaders have plenty of activity and little confidence about what's real, why deals stall, or where follow-up broke.

WHY IT MATTERS — More automation on top of bad targeting produces more of the wrong thing, faster.

05

Customer automation that costs trust when it fails

RISING

Bots are saving touches and creating churn risk wherever resolution or human handoff breaks down.

WHY IT MATTERS — Gartner found only 27% of customers would try a chatbot again after one bad experience.

The signal board.

↑ RISING → HOLDING ↓ COOLING • NEW

↑ RISING

The ROI question got an executive-grade proof point.

Gartner's September 1 release gave CFOs language they didn't have a month ago. Rising spend, 22% scaled, incomplete visibility, and a measurable behavioral difference between high and low performers. Expect that to show up in board decks this quarter.

• NEW

Customer patience on AI service now has a number.

Gartner's September 2 data: 27% would retry a chatbot after a bad experience, and 87% consider human access essential when GenAI is used in service. Their own recommendation is telling — make the bot a connector to humans, not a containment device.

↑ RISING

Small-business conditions got harder in a specific way.

NFIB's September 8 report showed actual sales deteriorating in August, the weakest reading since November 2025, with inflation climbing back to the top spot as owners' single biggest problem. Optimism held near its long-run average. That combination doesn't mean panic. It means selective spending, and it makes fuzzy-payback projects much harder to sell.

↑ RISING

Agent governance stopped being philosophical.

The conversation moved from “responsible AI” principles to identity, policy enforcement, observability, token cost, and control. That shift usually precedes procurement.

Where budget is actually moving.

Money is moving downstream. Away from model licenses and proofs of concept. Toward integration, governance, workflow redesign, implementation, measurement, and change management.

1 WHAT'S GETTING FUNDED

Hands-on implementation capacity, agent control and governance tooling, integration and data readiness work, CRM adoption and reporting fixes, and customer-journey redesign around escalation rather than deflection.

2 WHAT TRIGGERS THE SPEND

A failed pilot. A board asking why AI hasn't reached the P&L. An agent about to receive write permissions. A security or regulatory review. A CRM or contact-center renewal. A headcount freeze that creates a capacity mandate. Data cleanup suddenly getting executive sponsorship — that one is almost always downstream of an AI project that hit a wall.

3 WHO HOLDS THE BUDGET

CFO, CIO, COO, CRO, and in regulated deployments, the CISO. The operational champions are RevOps, enterprise architecture, service leadership, and transformation teams.

4 WHAT'S GETTING EASIER TO JUSTIFY

Anything with a named baseline, an owner, and a payback window. What's getting harder: anything described as transformation.





SIMMELLE'S TAKE

SIMMELLE ZANE | PARKWEST-CX

I don't think most companies have an AI measurement problem. I think they have a design problem that measurement is finally exposing.

You cannot calculate return on a workflow you never baselined. If nobody wrote down what the process cost, how long it took, how often it broke, and who owned the outcome — then the ROI conversation was always going to end in a shrug. The absence of numbers isn't a reporting gap. It's evidence the work was never examined before it was automated.

“Don't automate a process you wouldn't intentionally design today.”

Which brings me to the rule I keep coming back to: don't automate a process you wouldn't intentionally design today. A lot of AI spend has gone into making bad processes faster. That's not leverage. That's amplification.

The customer-service data is the sharpest example right now, and I'd argue it's being misread as a chatbot story. It isn't. Eighty-seven percent of customers wanting human access is not a verdict on the technology. It's a verdict on companies that deployed containment and called it service. Automate the repeatable. Keep humans where judgment, empathy, and exception authority actually live. The moment you engineer the escape hatch as an afterthought, you've built your churn machine.



Same logic on the revenue side, and here's the contradiction I find genuinely useful. Two-thirds of B2B buyers say they'd prefer a rep-free experience. Nearly seven in ten want a salesperson to validate what AI told them. Those numbers only look contradictory if you think sellers distribute information. They don't anymore. They validate, interpret, and reduce risk. If your sales model is still built around being the source of information, the buyer already worked around you before your CRM knew they existed.

And on agents: govern them like identities, not like features. What can it read, change, spend, approve, and send outside the building? Where must a person intervene? What gets logged? If you can't answer those before production, you haven't approved a use case. You've approved an exposure.

Find the leak before you buy the tool.

It is still, by a wide margin, the cheapest thing on the list.



SIMMELLE ZANE · PARKWEST-CX



Not the story yet. Possibly next.

Five signals that did not make this edition's main narrative but are worth watching.

01 **Agent-to-agent interaction**

Most governance thinking still assumes a human on one end. That assumption has a short shelf life.

02 **Data provenance and lineage**

Risk and compliance teams are starting to ask where model inputs originated. Very few companies can answer cleanly.

03 **Third-party AI vendor risk**

Executives are increasingly uneasy about how AI vendors handle sensitive data, security data in particular.

04 **Brand accuracy in AI search**

Not ranking — accuracy. What AI systems say about your company is becoming its own problem category.

05 **“Botsitting”**

Early evidence that agents create their own supervision and maintenance workload. Agent count is a terrible success metric.





For every AI initiative we're funding right now — can we name the baseline it's supposed to improve, the person accountable for the benefit, and the condition under which we'd shut it down?

If any of the three is missing, you don't have an investment. **You have a subscription.**

What sits underneath this edition.

- **Gartner** — AI scaling and spending research (Sept 1); customer service and GenAI research (Sept 2); B2B buyer behavior research
- **NFIB** — Small Business Economic Trends, August data released Sept 8
- **Federal Reserve Bank of Atlanta** — executive survey on AI, productivity, and workforce effects
- **Google Cloud / Accenture** — forward-deployed engineering capacity announcement
- **Boomi, SAP** — agent control plane, governance, and integration positioning
- **HubSpot** — CRM adoption and reporting programming
- Current first-person commentary from CFOs, revenue leaders, and technology practitioners



The Executive Frustration Brief

Published monthly by **PARKWEST**

PARKWEST helps businesses identify where technology and AI can create measurable business value before they invest in more tools.

If this was forwarded to you, visit parkwest-cx.com/executive-frustration-brief.

If you are wrestling with one of these issues, reply or message Simmelle.