



MAY 2026

THE 2026 STATE OF CONVERSATIONAL AI

RASA.COM

# The 2026 State of Conversational AI

Enterprise leaders share their thoughts on performance, confidence, and control as AI adoption grows

## EXECUTIVE SUMMARY

# Five notable findings from the research

01

## Enterprise AI programs are scaling faster than the confidence behind them.

Most surveyed enterprises (67%) are expanding or scaling their conversational AI solution, yet average confidence in AI's ability to handle complex conversations sits at just 4.37 out of 7. The gaps between momentum, conviction, and support make up the defining tension of this report.

02

## Control has replaced capability as the top concern.

Most surveyed leaders (60%) rank "black box" issues or compliance as their #1 challenge, ahead of integration, deployment complexity, and resource constraints. The question enterprises are wrestling with isn't whether AI is smart enough, but whether they can understand, govern, and stand behind what it does.

03

## Transparency in AI is key.

Nearly all respondents (93%) say AI transparency is "very important" or "critical," and nearly half (43%) say they won't deploy a solution without it. Enterprises are translating that demand into concrete architecture and infrastructure decisions. Two-thirds (66%) require on-premise or own-cloud deployment control, and nearly two-thirds prefer hybrid architectures that combine LLM flexibility with deterministic logic (63%) over fully agentic systems (13%).

04

## Getting AI live isn't the hard part.

Achieving performance metrics is the most commonly cited pain point along the AI journey, outpacing deployment by almost 3:1. Enterprises aren't asking for more powerful technology. They're asking for more realistic benchmarks, clearer paths to value, and simpler best practices for what "good" actually looks like.


05

## No two sectors or roles are fighting the same battle.


Where enterprises get stuck, what they're most worried about, and how they define success varies significantly by industry, function, and title. Generic, one-size-fits-all solutions and guidance rarely account for these differences. The organizations best positioned to succeed are the ones working with partners that understand their unique needs and constraints.

# Table of contents

- 01 [Methodology](#)
- 02 [Introduction](#)
- 03 [The State of Play](#)
- 04 [The Central Challenge](#)
- 05 [Building for Control](#)
- 06 [Where the Journey Breaks Down](#)
- 07 [Why Context Is Everything](#)
- 08 [Key Takeaways](#)
- 09 [Conclusion](#)



# 01 Methodology



METHODOLOGY

# About this report


Rasa surveyed 30 enterprise-level decision-makers across finance, healthcare, retail/e-commerce, government, and telecom in February 2026. Respondents included technical, product, and customer operations functions, with titles ranging from director level to the C-suite.

This report shares what they told us about their conversational AI programs, from how they're building to what they wish they'd known sooner.

Given the sample size, we treat these findings as directional signals rather than statistically significant data. The leaders we polled offer something just as valuable in a complex and fast-moving market: candid, firsthand perspectives from the people actively building and managing conversational AI systems at scale.



## 02 Introduction



## INTRODUCTION

# AI adoption is no longer the story. Getting it right is.

When it comes to conversational AI, most enterprises are past the stage of asking, “Should we build it?” and are now actively wrestling with, “How do we make it work?” This second question is much more difficult and far less glamorous to answer. How do you build a solution that holds up under compliance scrutiny, at scale, and across every channel your customers use?

To understand where the industry stands — not according to vendor roadmaps or analyst projections, but real practitioners actually building and managing these systems — we asked enterprise leaders what’s working, what isn’t, and what they’re still looking for. Their answers challenge several of the assumptions that have shaped conversational AI discourse in recent years. In many ways, the market is

more experienced, more pragmatic, and more wary of hype than the headlines suggest.


The numbers bear that out. The global conversational AI market is [projected to grow](#) at 23.7% CAGR to \$41.4 billion by 2030, led by chatbots as the most common application and retail as the most active adopter. A recent [McKinsey](#) study found that 88% of enterprises now regularly use generative AI for at least one business function, up 10 percentage points from a year before. And [Gartner](#) expects that as many as 40% of enterprises will have embedded AI agents into their applications by the end of this year. But momentum and maturity aren’t the same thing. A majority of enterprises using generative AI (62%) remain in the experimentation or pilot phase, and the organizations

scaling fastest are often outpacing the benchmarks, governance frameworks, and institutional knowledge they need to do AI well.

That’s the tension at the heart of this report. For many organizations, the gap between the pace of adoption and the level of available support is more than a growing pain. It’s the primary obstacle to realizing the value they’ve already committed to pursuing.

In the pages that follow, we explore why control has replaced capability as the top concern in conversational AI, where implementation hits real roadblocks, and what AI confidence looks like across different roles and industries. Our findings won’t necessarily tell you whether to build, but they will help you piece together what it actually takes to succeed once you do.

## 03 The State of Play



## THE STATE OF PLAY

# Most enterprises are scaling their AI programs, but confidence remains limited.

This isn't a market still debating whether AI belongs in customer conversations, but one figuring out how to make AI-driven conversations more effective, more consistent, and more defensible to leadership teams. Today's enterprise leaders are already managing live systems, iterating on outcomes, and making strategic decisions about what to build next.


**67%** of respondents are planning to expand or scale their conversational AI solution over the next 12 months

**20%** say they're skeptical that AI can reliably handle complex customer conversations in their industry, and only 3% have full confidence

In fact, the majority of respondents (67%) are actively expanding or scaling their conversational AI program. Very few (7%) are still planning their first deployment, and a small but significant number are pausing or reducing their AI initiatives (3%). Another 10% are exploring alternatives to their current solution, with the remaining enterprises taking no action this year — either because they're happy with their current solution (3%) or still determining their approach (10%).

This broadly tracks with a market that shows no signs of slowing down. According to [Gartner's 2026 CIO Agenda](#), 91% of enterprises are increasing their generative AI spending this year, with an average funding boost of 38%. The direction of the market is clear, and most organizations feel they can't afford to fall behind.

### Which best describes your organization's plans for conversational AI over the next 12 months?



Yet the same respondents who are expanding their AI programs display only middling levels of conviction. Across all industries, roles, and titles, confidence in the ability of conversational AI to handle complex customer conversations averages out to just 4.37 out of 7 — a barely passing grade.

Of the 20% of respondents who report they are “skeptical” of conversational AI’s abilities, all but one are scaling their programs anyway.

This isn’t so much a contradiction as it is a careful calculation. Enterprises aren’t expanding their AI programs because they’re sure they will work. They’re expanding because the economics of standing still are increasingly difficult to justify. Customer expectations for self-service quality [have risen](#)

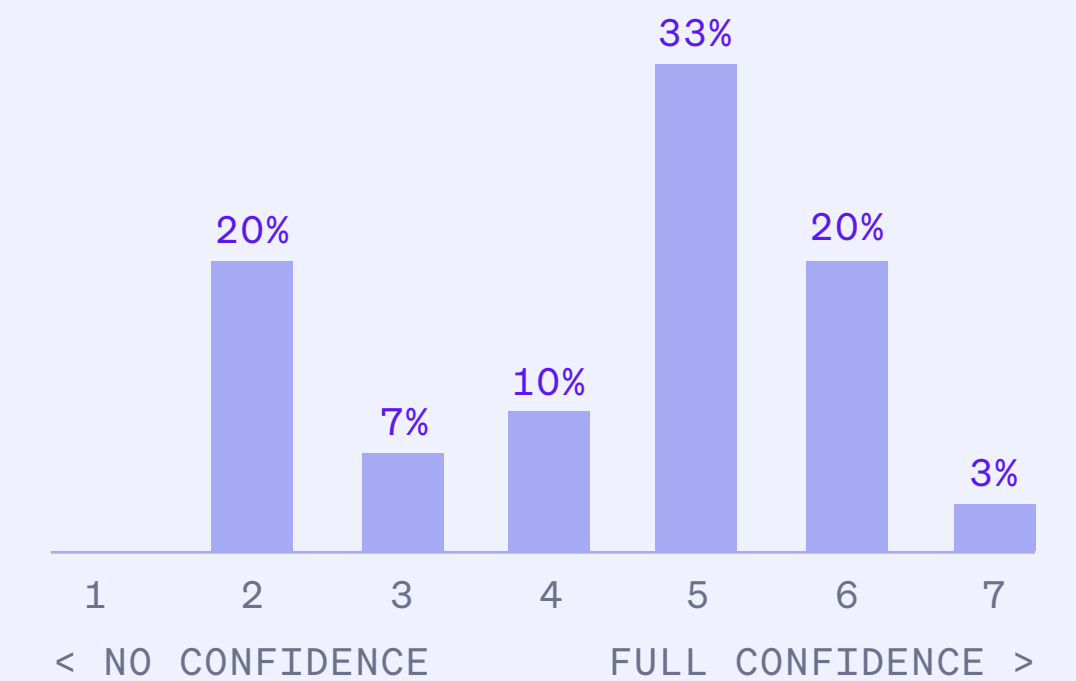
[dramatically](#). Meanwhile, competitors are deploying, and contact center costs aren’t going down. [Recent data](#) suggests that AI-handled customer interactions can cost a fraction of human-only ones and drive measurable gains in first-contact resolution, revenue enablement, and retention. Against that backdrop, the case for conversational AI is strong enough that lukewarm conviction has become an acceptable condition for moving forward, rather than a reason to wait.

Research also shows that bet isn’t always paying off, which explains why confidence is flagging. According to a widely cited [MIT study](#), despite tens of billions invested in generative AI, only a small fraction of initiatives advance past the pilot stage to deliver real returns. The organizations in our


survey understand this. But they also understand that the status quo carries its own cost. Most have already decided to scale. The question now is whether to build deliberately, with governance and measurement baked in, or reactively, once regulatory and competitive pressures force the issue.

That gap between scaling fast and scaling well is what shapes the rest of this report.

How confident are you that conversational AI can reliably handle complex customer conversations in your industry?



# 04 The Central Challenge



## THE CENTRAL CHALLENGE

# The most pressing problem isn't capability. It's trust.

Enterprises struggle most with whether they can understand, govern, and stand behind their conversational AI.

When asked to rank their main challenges, most surveyed leaders (60%) put “black box” problems — defined as a lack of visibility into AI behavior and output — or compliance concerns at the top of the list, ahead of integration, resource constraints, and deployment complexity.


**60%** of respondents identify “black box” issues or compliance as their #1 AI challenge

**93%** say transparency in AI decision-making is “very important” or “critical”

**43%** report they won't deploy a conversational AI solution without full explainability

That need for transparency is both widespread and resolute. Nearly all respondents (93%) say transparency is “very important” or “critical,” and almost half (43%) make it a hard requirement for deployment, meaning they won’t move forward without full explainability. All told, that suggests control has replaced capability as the foremost AI concern, and for many enterprises it isn’t negotiable.

**Please rank the following challenges in order of importance when it comes to implementing or scaling conversational AI at your organization.**



On the surface, control and compliance may seem like distinct issues. One is about understanding AI outputs (why did the AI model respond that way?), while the other is about accountability (how can we prove it responded the right way?). But they share a root cause, as both come down to visibility into what an AI system is doing and why. As one respondent put it, “If it can’t integrate with my systems, respect governance boundaries, and maintain reliability under pressure, it’s not a solution — it’s a liability.”

Intensity varies somewhat by function. Technical and IT respondents are universally uncompromising when it comes to transparency, with 100% rating

it “very important” or higher and 53% calling it “critical.” Product and innovation teams are close behind with a 90% high importance rating (40% “critical”). Meanwhile, customer operations teams, arguably the most focused on day-to-day performance outcomes, are slightly more pragmatic, with 80% rating transparency highly and 20% deeming it critical.

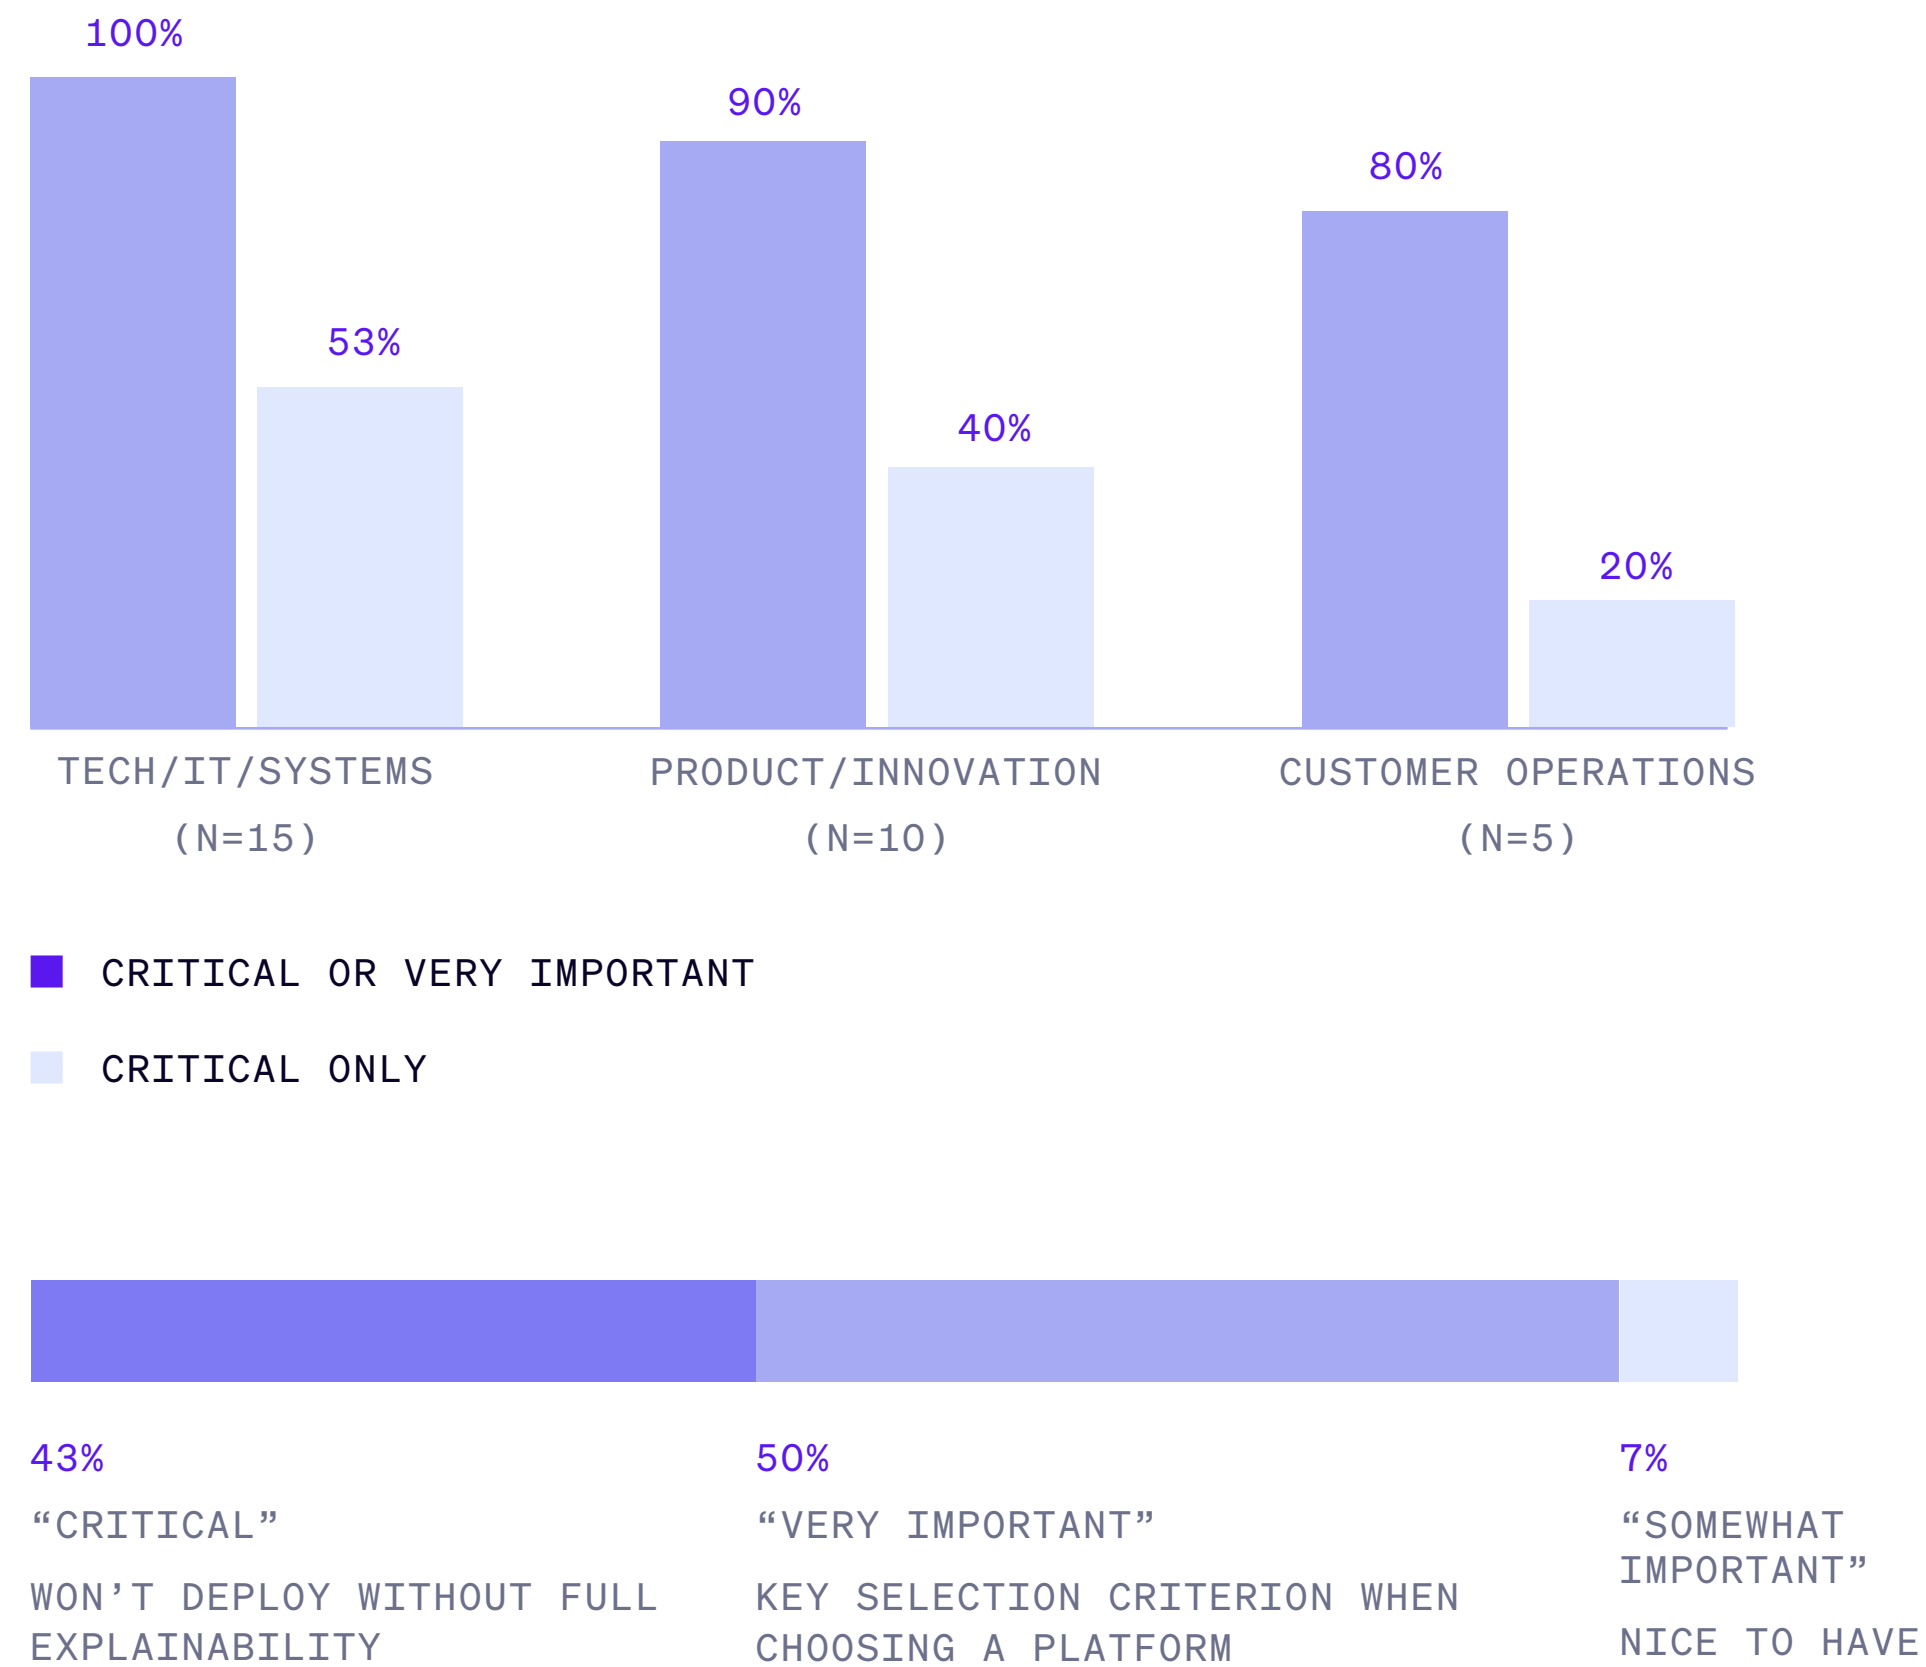
That demand for AI visibility doesn’t exist in a vacuum. Enterprises are navigating an increasingly complex and high-stakes regulatory landscape — from the EU AI Act’s transparency mandates to state-level AI enforcement in the U.S. — that gives their desire for control real legal and financial teeth.

“IF IT CAN’T INTEGRATE WITH MY SYSTEMS, RESPECT GOVERNANCE BOUNDARIES, AND MAINTAIN RELIABILITY UNDER PRESSURE, IT’S NOT A SOLUTION — IT’S A LIABILITY.”

**Rich**

Chief Information Technology Officer,  
Healthcare

How important is transparency in your conversational AI (i.e., that you can see and understand how your conversational AI model makes decisions)?



## The regulatory backdrop

The enterprise demand for AI transparency goes beyond a preference and carries regulatory implications across multiple jurisdictions.

In the United States, a patchwork of state-level AI laws and enforcement actions has created what legal analysts see as overlapping jurisdictional risk, where AI missteps can trigger investigations and lawsuits from multiple directions at once. AI-related issues contributed to over \$4 billion in FINRA and SEC penalties against financial services firms in a single year, driven by biased outputs, misleading communications, and inadequate supervision of automated systems.

In the EU, the AI Act mandates documented risk management, human oversight, logging, and transparency for high-risk AI systems across finance, healthcare, and public services, with fines up to €35 million or 7% of global annual turnover. The General Data Protection Regulation (GDPR) also requires that organizations provide “meaningful information” about the logic involved in automated decisions with significant effects.

For enterprises operating across regions, these overlapping frameworks make the question of how (and where) AI runs a first-order governance decision.

Considering this context, enterprise concern isn't surprising. What is surprising is where it doesn't show up.


Deployment complexity has long been treated as the primary obstacle in enterprise AI, but only 3% of respondents identified it as their top challenge. That doesn't necessarily mean deployment is easy, but it does suggest that the issues that emerge once a system is live are harder and more persistent than getting it up and running in the first place.

Equally notable is that internal resistance ranked eighth out of eight challenges, with only 7% listing it as their top concern and nearly half (47%) placing it dead last. This inverts a common assumption in enterprise tech, where securing organizational buy-in is considered a critical roadblock to large-scale initiatives. With conversational AI, the blockers appear to be more

structural — technical architecture, governance frameworks, and measurement systems — than organizational. For anyone scoping a conversational AI implementation, it's a useful signal that investment in change management may be better directed toward the technical and performance areas that are actually slowing these programs down.



# 05 Building for Control



## BUILDING FOR CONTROL

# Enterprises are reframing control as a design decision.

When enterprises say they want control over their conversational AI solution, they don't just mean they want to understand what it's doing. They also mean they want to govern where it runs, how it's designed, and what it's allowed to do autonomously.

In other words, control isn't a single requirement with a single solution. It shows up differently at different levels, from infrastructure decisions about where AI runs to architectural choices about how much autonomy it's given. Our findings reflect that in two distinct but complementary ways.

**66%** of respondents consider on-premise or own-cloud deployment to be "very important" or "essential"

**63%** prefer hybrid architecture that combines LLM flexibility with deterministic logic

**13%** opt for fully agentic systems

## Where AI runs

Two-thirds of surveyed enterprise leaders (66%) consider on-premise or own-cloud deployment either “very important” or “essential” when it comes to conversational AI. Half of those leaders are driven by regulated or security mandates. The other half simply prefer it that way.

That means deployment control has become a business demand in its own right, independent of formal compliance obligations. As one VP of IT put it, “I need real operational governance: versioning, approvals, role-based admin, dev/test/prod separation, and change logs so IT can control behavior like a production system.”

Only 17% of respondents say cloud-based deployment is fine, making generic shared-cloud environments a distinct minority preference among surveyed enterprises.

This largely aligns with broader industry trends. Forecasts from [IDC](#) suggest that, by 2028, roughly 75% of enterprise AI workloads will run on hybrid, fit-for-purpose infrastructure, with spending on dedicated single-tenant cloud environments growing [more than twice as fast](#) as shared cloud. At the same time, the EU AI Act and GDPR are making data residency and sovereign processing [central compliance considerations](#) for high-risk AI systems operating in Europe, turning the question of where AI runs into a first-order governance decision rather than an afterthought.

“I NEED REAL OPERATIONAL GOVERNANCE: VERSIONING, APPROVALS, ROLE-BASED ADMIN, DEV/TEST/PROD SEPARATION, AND CHANGE LOGS SO IT CAN CONTROL BEHAVIOR LIKE A PRODUCTION SYSTEM.”

**Michelle**

Vice President of Information Technology,  
Healthcare


## How AI is designed

The same instinct shows up in how enterprises design their AI. When asked about their ideal conversational AI architecture, most respondents (63%) chose hybrid models that combine LLM flexibility with deterministic logic, far outpacing fully agentic systems (13%) and deterministic-first approaches (10%). Another 13% are still assessing their options.

It's tempting to read the preference for hybrid architecture as fence-sitting, but this would be a mistake. Enterprises are making a considered judgment about where each approach adds value. LLMs are exceptional at understanding intent and generating natural responses.

Deterministic logic is better suited to parts of a conversation where variance isn't acceptable, like billing disputes, regulatory disclosures, multi-step authentication flows, and high-stakes actions where getting it wrong carries real consequences.

With this in mind, the preference for hybrid doesn't reflect indecision. Respondents are making a deliberate judgment call about which approach suits which part of the conversation.



Which best describes your ideal conversational AI architecture?

- FULLY AGENTIC ■
- PREFER HYBRID ■
- STILL EVALUATING ■
- DETERMINISTIC-FIRST ■

That preference extends beyond architecture into service design. In open-ended responses, enterprise leaders consistently described plans for escalating seamlessly to human agents when conversations become too sensitive or too complex, passing along the context, transcript, and reason for handoff.

This suggests they're not designing toward full AI replacement of human service, but toward hybrid models from day one, with AI handling volume and humans handling judgment — an arrangement [McKinsey](#) finds has the most overall impact in customer-facing and regulated environments.

The 13% of respondents who prefer fully agentic architecture deserve a closer look. They're not outliers and may in fact be early indicators of where the market is heading.

This cohort is simultaneously the most architecturally ambitious and the most governance-aware in our survey. A majority (75%) ranked control or compliance as their top challenge, meaning they're not naive about the risks agentic models can introduce. They're building despite them. That pattern is more consistent with early adopters who have a clearer view of the governance work required than with organizations underestimating what's involved.


Their experience surfaces a tension the rest of the market will eventually face. A [2026 Deloitte survey](#) found that, while 74% of companies plan to deploy agentic AI within two years, only 21% have a mature governance model for autonomous agents in place. The gap between ambition and readiness isn't unique to the agentic cohort and is actually a market-wide condition. The difference is that these teams are confronting it now, in production, rather than in theory.

The governance tooling, benchmarking infrastructure, and organizational readiness required to operate fully agentic systems at scale are still maturing. For most enterprises today, hybrid architecture offers the best available balance of flexibility and control. But as those supporting capabilities develop, the calculus may shift, and the organizations already learning what agentic governance requires in practice will have a meaningful head start.

For now, market [narratives](#) around agentic AI have been loud and confident, but [actual deployment lags behind the rhetoric](#). The enterprises actually building these systems are largely taking a measured view — not because they're behind the curve, but because they understand what's at stake when AI gets it wrong.



## 06 Where the Journey Breaks Down



WHERE THE JOURNEY BREAKS DOWN

# The hardest part isn't getting AI live. It's knowing whether it's working.

Getting a conversational AI system up and running is hard, but it's a solvable problem with a defined finish line. What comes after is more difficult and less supported.

In our survey, measuring and improving performance was the single most commonly cited hardest phase of the AI journey, outpacing deployment (10%) by a margin of nearly three to one.

This matters because it pinpoints where and when enterprises actually need help. Planning (20%) and building or securing budget (17% each) account for the next biggest pain points, suggesting early-stage stumbles are also common before deployment friction even enters the picture. Regardless, different organizations are hitting walls at different stages, and one-size-fits-all guides and playbooks can't address what most enterprises are actually struggling with.


**27%** of respondents say achieving performance metrics is the hardest phase in the AI journey, with only 10% listing deployment

**60%** say clearer paths to success and faster time to value would make implementation easier

When asked what would help most, responses were notably operational rather than technical. Stronger tooling and more powerful AI features didn't rank among practitioners' top requests.

Instead, they're asking for greater clarity, including clearer paths to value (60%), clearer benchmarking data (53%), clearer best practices for implementation (50%), and more hands-on guidance during deployment (50%). The top two requests are essentially variations of the same ask: give us something concrete to measure against, and show us what "good" actually looks like in practice.

### What phase of your conversational AI journey has been the most challenging?



What would make implementing and scaling a conversational AI solution significantly easier for your team?

- 60% want clearer paths and faster time to value (18/30)
- 53% want clearer benchmarking data (16/30)
- 50% want clearer best practices for implementation (15/30)
- 50% want more hands-on guidance during deployment (15/30)
- 47% want pre-built templates or use case examples (14/30)

The gap between scaling and knowing whether you're succeeding is reflected in the wider market. [Studies](#) show that, even as AI models become more capable, benchmarking methods remain biased, narrow, and highly sensitive to small evaluation changes, leaving enterprises with limited and inconsistent tools for judging real-world performance.

As a result, many organizations are forced to piece together usage, containment, and resolution data from a "patchwork" of logs, dashboards, and manual analysis, with no universal analytics stack to streamline the process.

This is an important gap. [BCG's 10-20-70 framework](#) argues that only around 30% of successful AI deployment comes down to

algorithms and infrastructure. The remaining 70% is about people and process — the governance frameworks, measurement systems, and operational disciplines where respondents say they're undersupported.

In fact, when asked openly what they most need from a conversational AI solution to be successful, the terms respondents used most frequently weren't about features or capabilities. They were operational: accuracy, governance, control, integration, compliance, and reliability. This is the language of practitioners who have been managing AI systems long enough to know exactly where and when they break down.

**What are the most important capabilities or qualities you need from a conversational AI solution to be successful?**

The picture becomes sharper when enterprises say what they’re actually trying to measure. Response accuracy tops the list of critical or very important metrics at 90% — a full ten percentage points above compliance adherence (80%), customer satisfaction (77%), and productivity improvements (70%).


It’s a clear demand for better ways to verify that AI models are performing correctly, which requires evaluation frameworks, testing infrastructure, and review processes that most enterprises are still building out.

The gap between accuracy and automation rate (63%) is particularly telling. While maximizing the volume of conversations AI can handle is what dominates vendor marketing, what enterprises are primarily trying to ensure is that the conversations it does handle are managed correctly.

**How important is each of the following metrics to you when it comes to measuring the success of conversational AI?**

METRIC	CRITICAL OR VERY IMPORTANT
RESPONSE ACCURACY	90% (27/30)
COMPLIANCE ADHERENCE	80% (24/30)
CUSTOMER SATISFACTION (CSAT/NPS)	77% (23/30)
AGENT PRODUCTIVITY IMPROVEMENT	70% (21/30)
AUTOMATION RATE	63% (19/30)
CONTAINMENT/DEFLECTION RATE	53% (16/30)
COST PER CONVERSATION	53% (16/30)
FIRST CALL RESOLUTION (FCR)	53% (16/30)
AVERAGE HANDLING TIME (AHT)	33% (10/30)

# 07 Why Context is Everything



## WHY CONTEXT IS EVERYTHING

# No two sectors or roles are fighting the same battle.

The aggregated findings in this report are powerful, but there's real value in specifics. Enterprises in different industries are hitting different walls at different stages of the AI journey for different reasons. And within those organizations, different roles are navigating the same challenges in ways that don't always follow expected patterns.

Generic guidelines and best practices rarely account for these distinctions, which is likely why so many enterprises feel undersupported. Let's dig into some of the nuances.

## Variations across industries

### Financial services: the least confidence, the most compliance intensity

Every financial services respondent rated transparency as either “critical” or “very important,” with half saying it’s a make-or-break factor for deployment.

Compliance tops their list of challenges at nearly four times the rate of their healthcare counterparts. Meanwhile, their average confidence score (4.12/7) is the lowest of any industry in our survey.

This combination of high compliance pressure and low confidence tells

a clear story. Financial services organizations are typically past planning and building and have cleared the early-stage hurdles. They’re now fighting the more persistent challenge of proving their AI systems are working in an environment where “close enough” isn’t an acceptable answer. Performance measurement is where these teams struggle most, which tracks with the stringent accuracy and auditability standards of a sector where AI-related compliance failures can carry hefty consequences.

**100%** of financial services respondents rate transparency as “very important” or “critical”

**38%** cite compliance as their top challenge

**4.2%** average confidence score is the lowest of any industry



## Healthcare: stuck earlier, but surprisingly confident

Healthcare offers a striking contrast to financial services. Compliance anxiety is surprisingly low given the heavily regulated environment, with only 11% of respondents citing it as their top concern. This could reflect a focus on administrative, non-clinical use cases where regulatory pressure is lighter, or it could signal these organizations are already accustomed to navigating strict oversight and have built compliance deep into their workflows. Either way, the majority of healthcare respondents (67%) are getting stuck earlier, at the building or budgeting

phase of their AI journey, suggesting they haven't yet encountered the compliance friction that overshadows some of their peers.

The relatively high confidence score (4.78/7) may reflect that optimism, or it could signal that the hardest challenges are yet to arrive. Healthcare leaders who are currently stuck at the building phase may find compliance and performance measurement pressures more acute once they reach deployment and scale.

**67%** of healthcare respondents encounter hurdles at the building or budgeting phase

**11%** cite compliance as their top challenge

**4.78** average confidence score is the highest of any industry



## Retail: more control-conscious than anyone would expect

Retail is the industry most would assume has the greatest flexibility around AI deployment models, as it's consumer-facing, fast-moving, and less encumbered by the regulatory frameworks that define financial services and healthcare. Our survey says otherwise.

Retail respondents require on-premise or own-cloud deployment at a higher rate than any other industry in our survey: 89%, compared to 75% in financial services and 44% in

healthcare. A possible explanation is that retail AI programs operate at significant scale, handling sensitive customer data and carrying direct reputational risk. A conversational AI system that underperforms in retail affects the way customers feel about (and whether they buy from) a preferred brand. Retail leaders appear to have internalized data sovereignty and performance reliability as core business requirements, even where formal regulatory mandates are fewer.

**89%** of retail respondents say on-premise or own-cloud deployment is “very important” or “essential,” with 44% reporting it's mandated by regulatory or security standards

**33%** list deployment as the most challenging phase in the AI journey



## Government and public sector: compliance first, everything else second

The government sample in our survey was comparatively small, but the pattern is clear enough to be instructive. Public-sector AI programs carry a fundamentally different kind of accountability than their private-sector counterparts. More than just a preference or a competitive advantage, transparency is often a civic demand. Citizens interacting with government AI systems have a reasonable expectation that those systems operate fairly, consistently, and in strict accordance with published policy. Every response, decision, and escalation is potentially subject to

public scrutiny in a way that doesn't apply anywhere else.

For government and public sector leaders, this means the stakes of getting AI governance wrong are both qualitatively different and quantitatively higher. Compliance topped their challenge list, and 100% rated transparency as "very important" or "critical," matching financial services as the most transparency-conscious industry in our survey — but for reasons rooted in democratic accountability rather than regulatory penalty.

INDUSTRY	AVG CONFIDENCE (/7)	TRANSPARENCY RATING	DEPLOYMENT CONTROL RATING	TOP CITED CHALLENGE
FINANCIAL SERVICES	4.12	100%	75%	Compliance & black box (tied)
HEALTHCARE	4.78	89%	44%	Black box
RETAIL/ E-COMMERCE	4.44	89%	89%	Compliance
GOVERNMENT/ PUBLIC SCTOR	4.33	100%	33%	Compliance

## Variations across roles and titles

### Confidence is uniformly modest across the organizational chart

One of the more surprising findings in our survey is what doesn't vary by seniority or function: confidence. Across roles and titles, confidence in conversational AI's ability to handle complex customer conversations lands within 0.25 points of the overall average. The common assumption is that executives tend to be more bullish on AI than those implementing it — in other words, that optimism decreases the closer you get to the technical reality. Our data doesn't support that, as moderate confidence is pervasive

from the boardroom down to the build team.

One exception worth noting is the range of responses within the C-suite, where confidence scores span from 2 to 7 — the widest variance of any group. While the average is broadly in line with other roles, that spread suggests disagreement at the leadership level. Not all executives are on the same page, and in some organizations, the gap between the most and least confident voices in the room may be consequential.

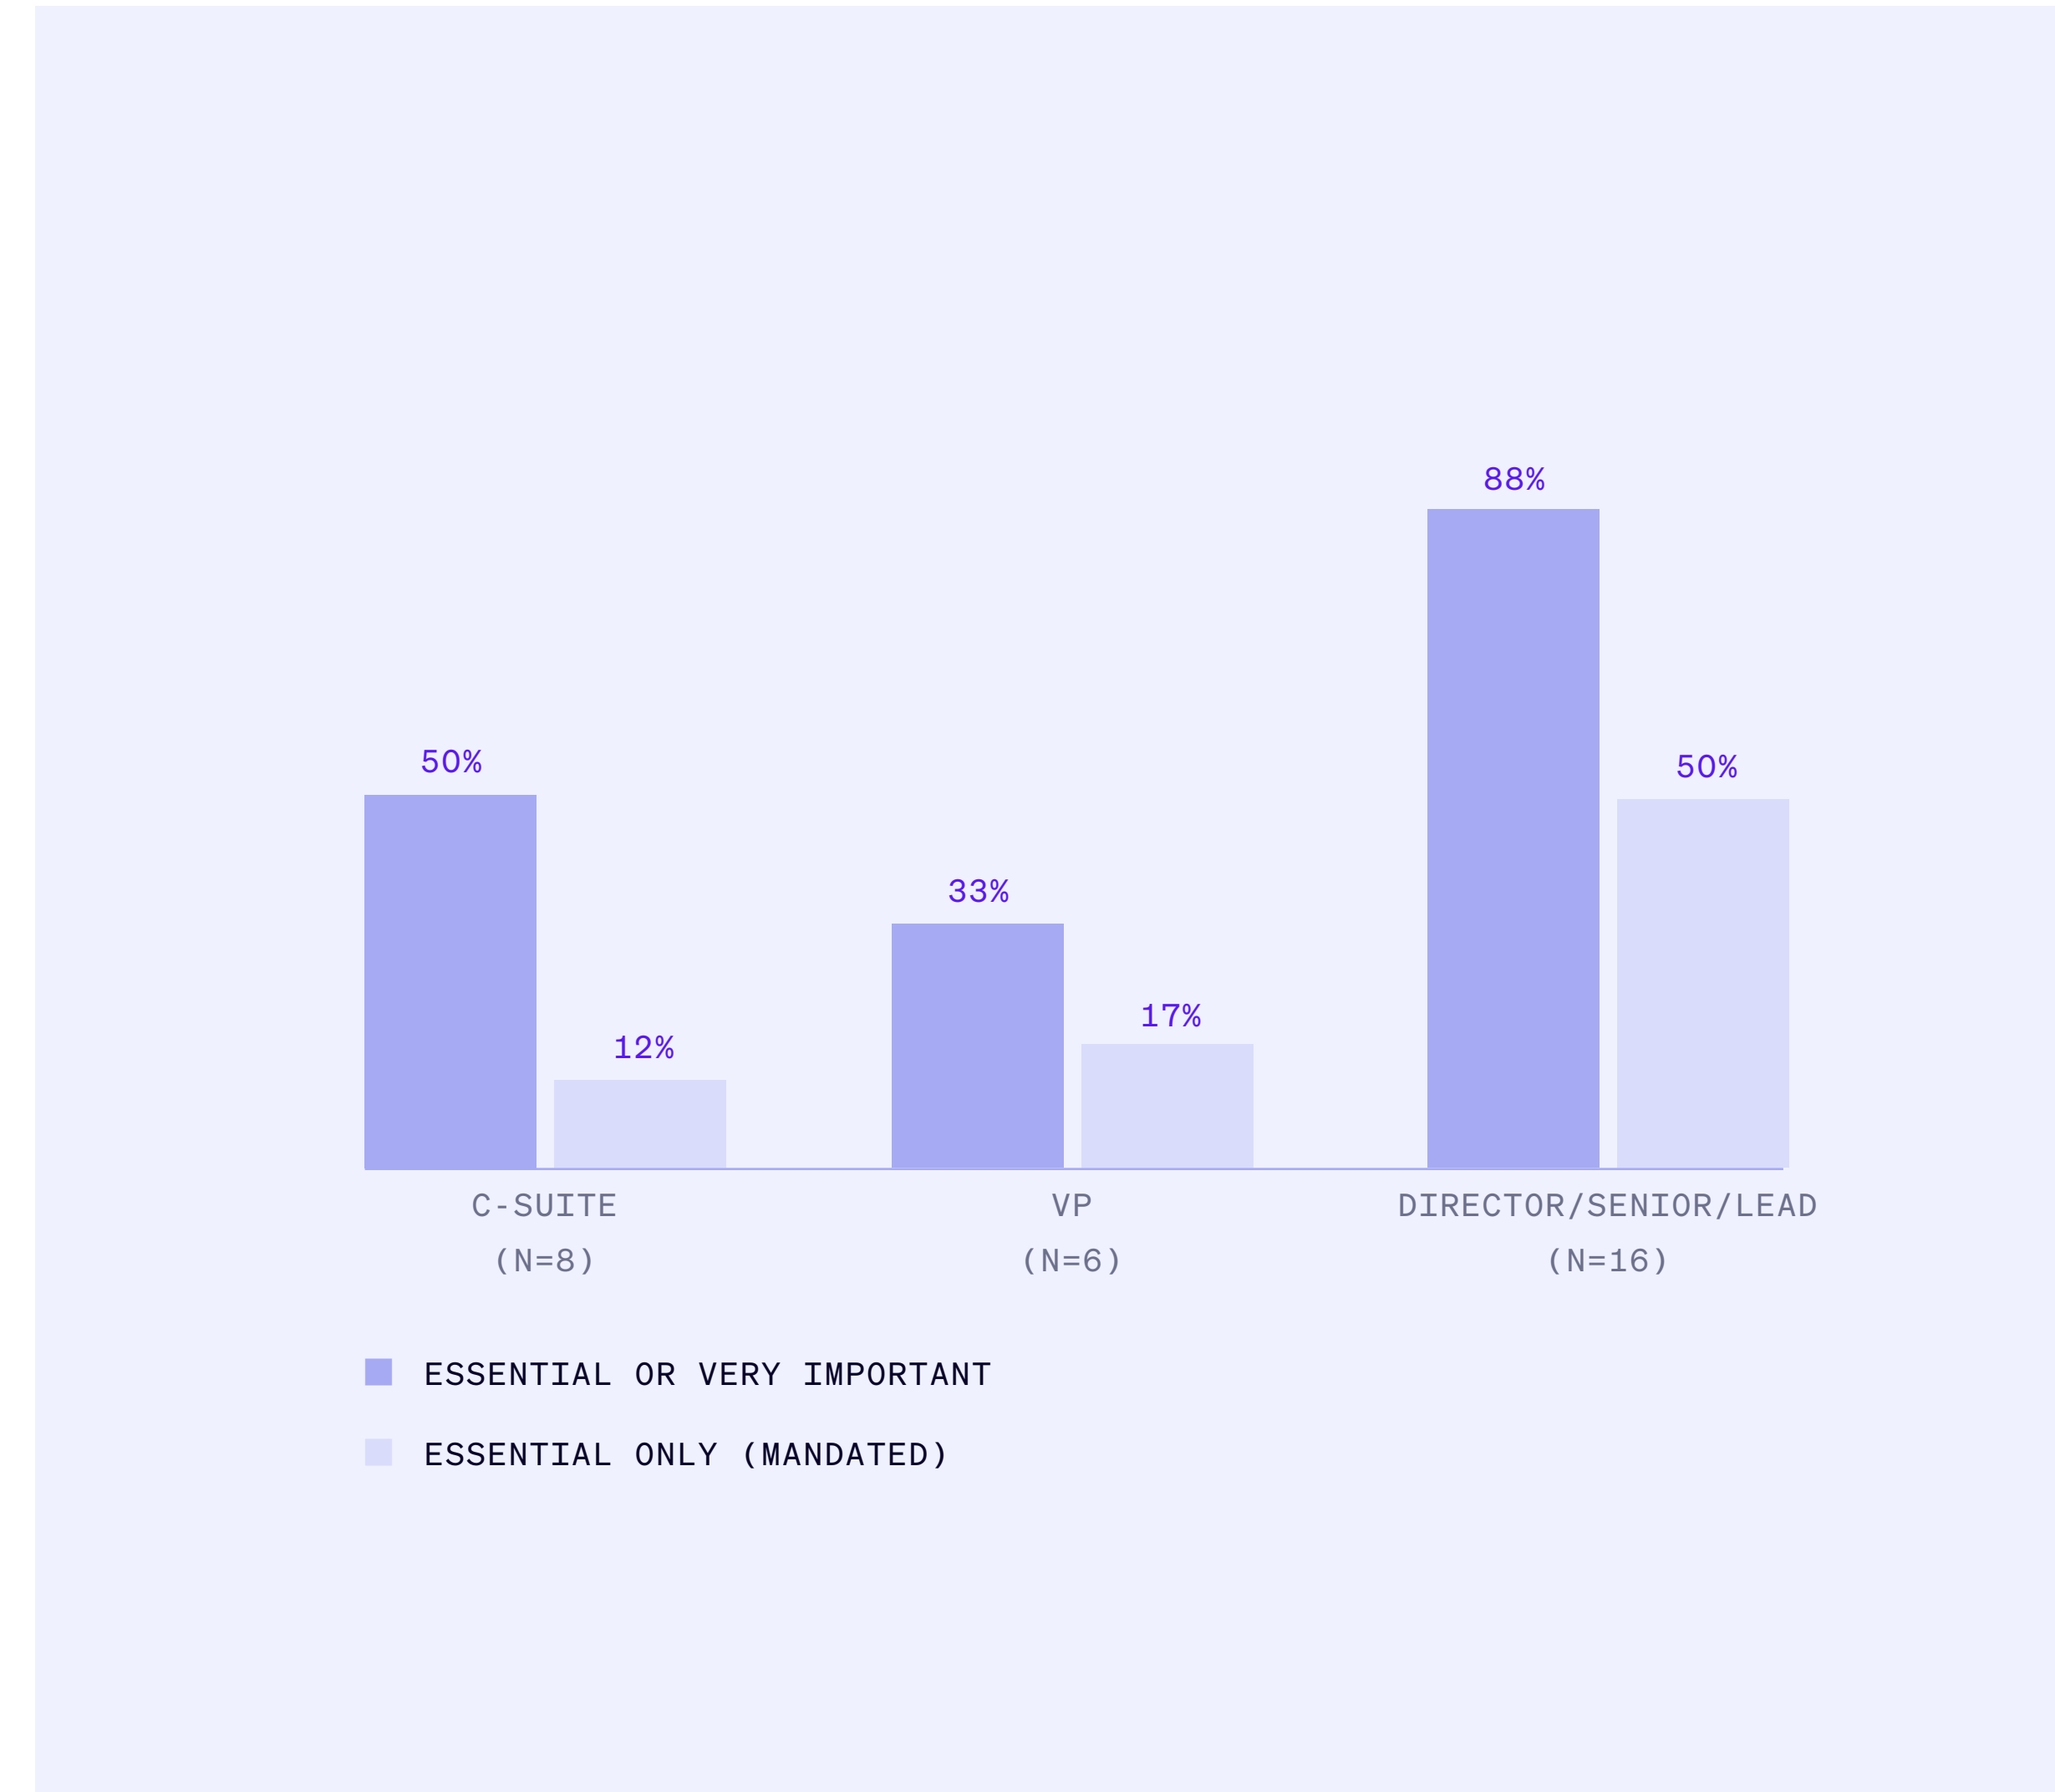
GROUP	N	AVG CONFIDENCE (Q9, 1-7 SCALE)
C-SUITE	8	4.25/7
VP/DEPARTMENT HEAD	6	4.50/7
DIRECTOR/SENIOR LEAD	16	4.38/7
TECH/IT/SYSTEMS	15	4.40/7
PRODUCT/INNOVATION	10	4.30/7
CUSTOMER OPERATIONS	5	4.40/7
OVERALL	30	4.37/7

## Product directors are driving deployment control more than engineers

The most counterintuitive finding here is about who is actually pushing for deployment control. Among our participants, 88% of directors and senior leads require on-premise deployment control, versus 50% of the C-suite and 33% of VPs. When you look at where that concern is coming from functionally, it isn't the technical teams. It's the product and innovation leaders.

Product and innovation roles show stronger deployment control preferences than technical roles, with 80% deeming on-premise or own-cloud deployment "very important"

and 60% deeming it "essential." For tech roles, those numbers are 60% and 13%, respectively. The implication is that deployment control isn't exclusively or even primarily a technical architecture concern. It's also a business architecture concern, driven by the people responsible for defining system requirements and scoping what gets built. If your implementation leads are asking for deployment control and your leadership isn't accounting for it in scoping decisions, that gap is worth closing before the next project begins.




## VPs are uniquely focused on integration

VPs are the only leadership group in our survey where legacy system integration tops the challenge list, above black box and compliance concerns. This likely reflects their structural position. VPs sit between strategic direction and technical execution, which puts them at the layer where integration friction is

most evident. They're close enough to implementation to feel the pain of connecting conversational AI to CRMs, contact center platforms, and legacy data sources — but far enough from the technical details to experience it as a blocker to business outcomes, rather than a solvable engineering problem.

ROLE GROUP	N	DEPLOY: ESSENTIAL OR VERY IMPORTANT	DEPLOY: ESSENTIAL ONLY
TECH/IT/SYSTEMS	15	60% (9/15)	13% (2/15)
PRODUCT/INNOVATION	10	80% (8/10)	60% (6/10)
CUSTOMER OPERATIONS	5	60% (3/5)	40% (2/5)

# 08 Key Takeaways



## KEY TAKEAWAYS

# Five principles for building conversational AI that scales

## 01 Solve for trust before you solve for capability.

The top challenge in conversational AI may not be whether the model is smart enough, but whether you can understand, govern, and stand behind what it does. Architecture decisions made early on around transparency, auditability, and deployment have longer-lasting consequences than feature decisions made later. The foundation matters more than the functionality built on top of it.

## 02 Define what “good” looks like before you scale.

More respondents cite achieving performance metrics as their most challenging phase — harder than deployment, building, and securing budget. If you don’t have a clear metrics framework before you scale, you won’t be able to identify or fix what isn’t working. Response accuracy and compliance adherence lead the list of what enterprises care most about measuring, and both require deliberate evaluation infrastructure to track reliably.

## 03 Know which battle your industry is actually fighting.

Financial services organizations are navigating compliance pressures and measurement issues in an environment with little tolerance for error. Meanwhile, healthcare teams are getting stuck at the building and budgeting phases, and retail is managing deployment control requirements that most observers would assume are looser than they actually are. Asking the right questions early on will help you ensure you’re solving for the right problems.

## 04 Listen to the people closest to implementation.


Directors and senior leads are more likely to require deployment control than the C-suite, and product and innovation teams are pushing for infrastructure governance just as hard (if not harder) than technical ones. The people defining system requirements and scoping what gets built often have a clearer view of the risks than the people approving the budget. If there's a gap between what your implementation leads are asking for and what leadership is accounting for, close it.

## 05 Hybrid architecture is a considered position, not a cautious one.

Most respondents prefer hybrid architectures that combine LLM flexibility with deterministic logic, and they're not hedging. They're making a specific, informed judgment about where each approach adds value. LLMs are adept at understanding intent, while deterministic logic is better at handling sensitive parts of a conversation where variance is unacceptable.

The agentic path isn't necessarily faster or simpler. It just moves the governance problem to a different part of the stack, one where tooling, benchmarks, and organizational readiness are still catching up. For most enterprises today, hybrid offers a more manageable path to production-grade control, though that balance will evolve as the ecosystem matures.

# 09 Conclusion




## CONCLUSION

# The gap between scaling and succeeding

Enterprise conversational AI in 2026 is a story of momentum and uncertainty running side by side. The enterprises in our survey are scaling, investing, and well past the point of experimentation. Most are doing it with only moderate confidence, unclear benchmarks, and a growing sense that the guidance and support available to them isn't keeping pace with what they're being asked to build and deliver.

This gap isn't a reason to slow down. The competitive and operational case for conversational AI remains strong, and the risk of falling behind too real. But it is a reason to build deliberately, with governance, measurement, and transparency frameworks treated as design imperatives from day one, rather than features to be tacked on later.

What the enterprise leaders in this survey make clear is that winning in conversational AI is less about what you build and more about whether you can stand behind it, understand it, measure it, and improve it with confidence. That's a harder problem than capability alone. And it's the one that matters most right now.



ABOUT RASA

# Built for enterprises that need more than \*off-the-shelf\*

Rasa is the developer platform for enterprise AI agents. Organizations in banking, healthcare, telecom, and retail use Rasa to deploy AI agents across voice and chat — on their own infrastructure, with full control over data, security, and behavior.

Trusted by Fortune 500 companies, Rasa ensures data privacy, security, and scalability for enterprises. Rasa is privately held, with funding from Accel, Andreessen Horowitz, and Basis Set Ventures.

LEARN MORE AT [RASA.COM](https://rasa.com)