RESEARCH REPORT

AI at Work

Where AI is delivering, and where it's quietly falling short in 2026

Leebry
by MacPaw

Intro

The "AI at Work" report starts with a simple question:
where is AI actually delivering value, and where is it falling short?

Adoption is accelerating, but most organizations are still working through the realities of implementation. Leadership expectations are rising, ownership is often unclear, and few teams have reliable ways to measure success.

At the same time, AI is delivering results, just not at the scale many expected. The most effective use cases are narrow, team-specific, and closely managed, while broader initiatives struggle to produce consistent outcomes.

Across every dimension in this research, one pattern stands out: the limiting factor is rarely the technology itself. It is the surrounding conditions, including data quality, governance, and how AI is introduced into the organization.

This report explores where those gaps exist, and what separates the organizations seeing results from those that are not.

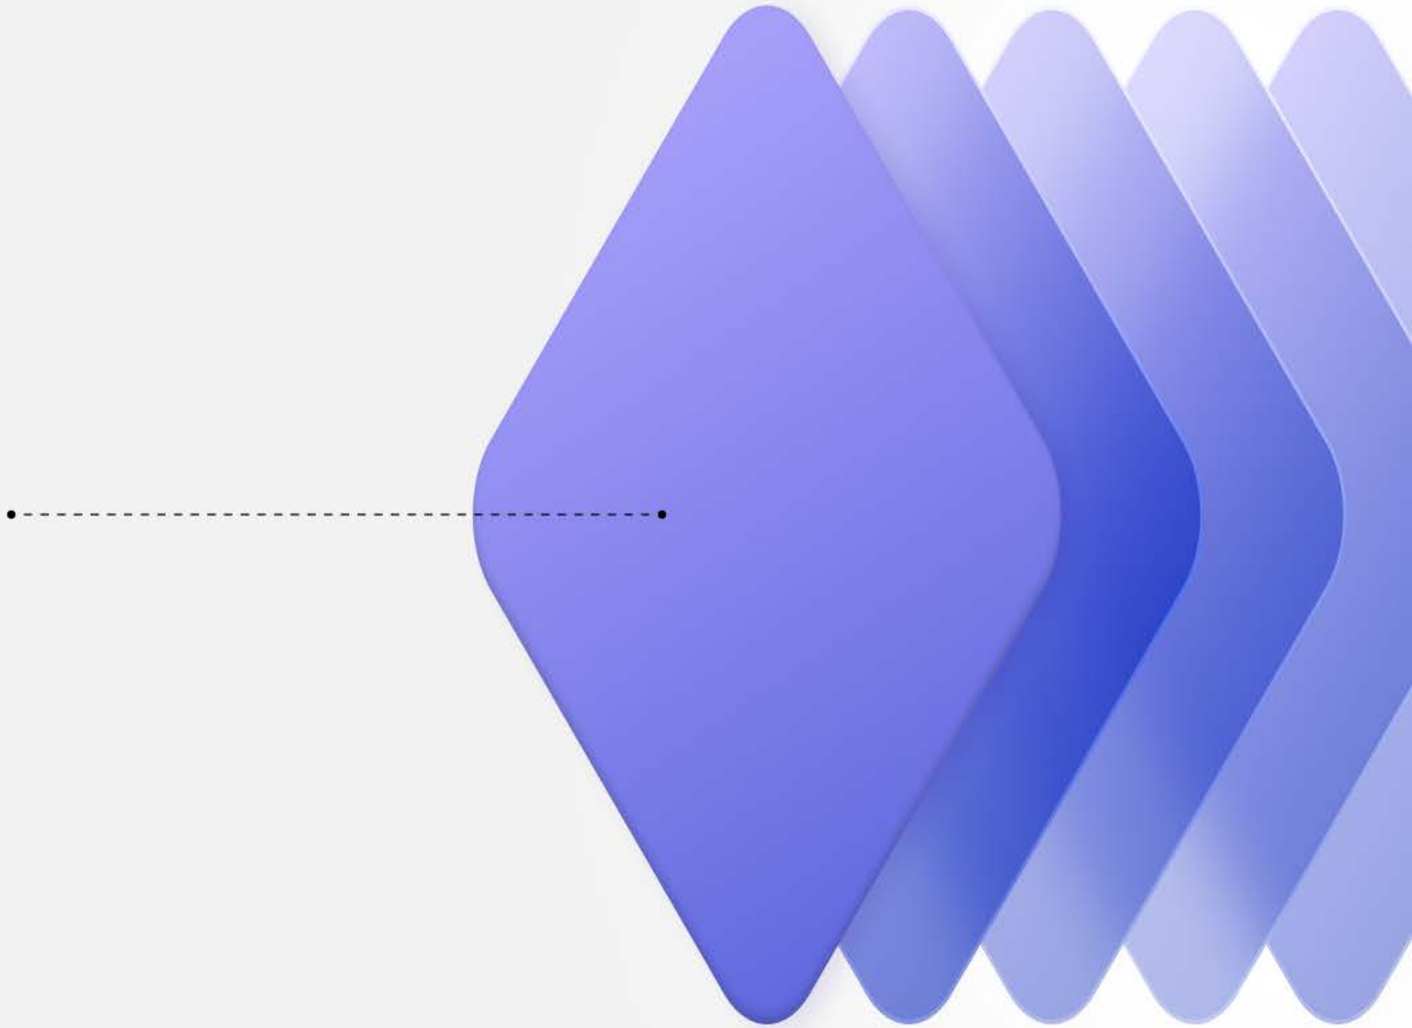

Contents

Key Takeways

Leadership expectations are outpacing execution.

61% of IT leaders report a gap between what leadership expects and what their teams can realistically deliver. That gap has consequences: organizations with strong alignment are 4.5x more likely to keep initiatives on track. Misaligned ones are quietly scaling back or stopping entirely.

Organizations are investing in AI they can't measure.

73% Despite continued investment, 73% of organizations lack formal, outcome-based ROI metrics. Only 17% say most of their initiatives have delivered measurable results. Without clear measurement, underperforming projects don't fail fast; they linger in uncertainty.

Data readiness is a hidden blocker.

90% of organizations deployed AI without auditing the internal knowledge it would access. 32% estimate their knowledge base is less than 50% accurate. Bad inputs produce bad outputs, regardless of how advanced the model is.

AI is not operating autonomously, it's being double-checked.

98% of organizations verify outputs before acting on them, and only 2% trust AI without human review. With limited visibility and real security risks, most organizations are still building that trust.

How This Research Was Conducted

This research was conducted in March 2026 via Wynter.
All 100 participants hold Director-level or higher IT roles.
Responses were collected across a wide range of sizes and industries across North America.

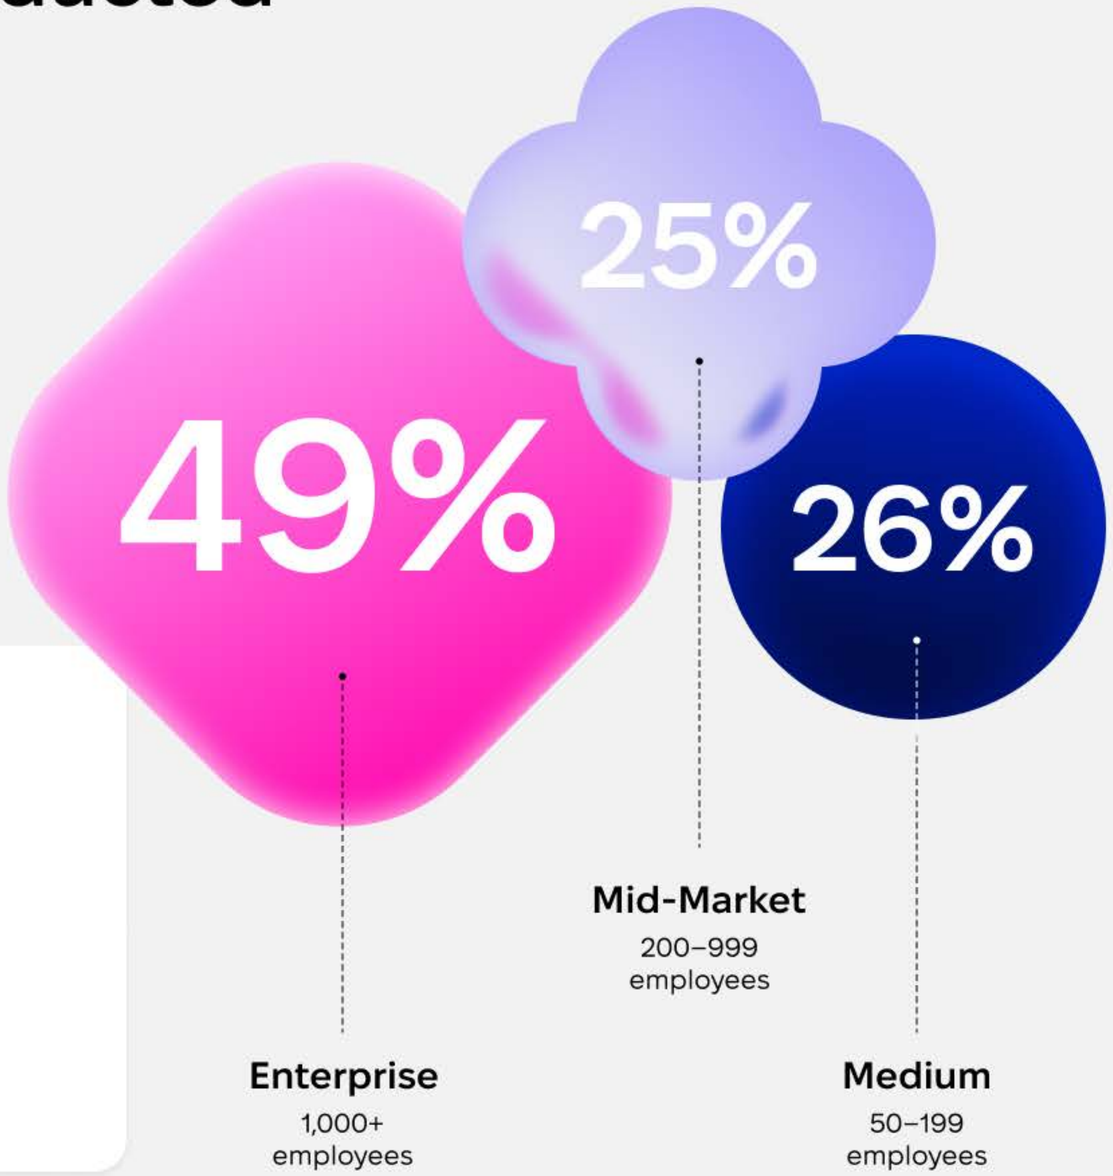

Industries represented include:
SaaS/software, professional and technical services, IT consulting, manufacturing, financial services, non-profits, healthcare, cybersecurity, automotive, government, and retail, among others.

* Methodological Note:
Analyses by company size is based on limited sample sizes. Percentage differences between groups should be interpreted with caution, as they are indicative rather than statistically conclusive. Readers are encouraged to focus on directional trends rather than precise percentage differences between groups

The Reality Gap in AI

AI adoption is accelerating, but expectations are moving faster than execution: 61% of IT leaders report a gap between what leadership expects, and what teams can realistically deliver today.

How aligned is leadership'sAI expectations with what your team can deliver today?

The gap looks different by company size, as the pressure is not evenly distributed.

Enterprise organizations

feel the most pressure: 65% report a moderate or significant gap, and only one respondent reported 2% say leadership has no specific expectations.

Mid-market organizations

are close behind at 64%, but 20% say leadership hasn't set expectations yet, leaving IT to define direction on their own

Medium companies

report the lowest expectation gap at 50%, but 27% say leadership has no AI expectations at all, meaning teams are adopting AI without a mandate or a framework.

Key Finding: Alignment Predicts Outcomes

3X

Companies that are misaligned on AI expectations are 3x more likely to stop or scale back initiatives. 33% of initiatives get affected in misaligned orgs. vs. only 15% in well-aligned organizations.

For the full breakdown of where AI initiatives are falling short, see page 13

Why AI ROI Remains a Black Box

Most organizations still do not have a clear way to measure whether AI is working.

73% lack formal, outcome-based metrics, relying instead on adoption rates or subjective feedback rather than clear business impact

How does your organization measure the ROI of AI initiatives?

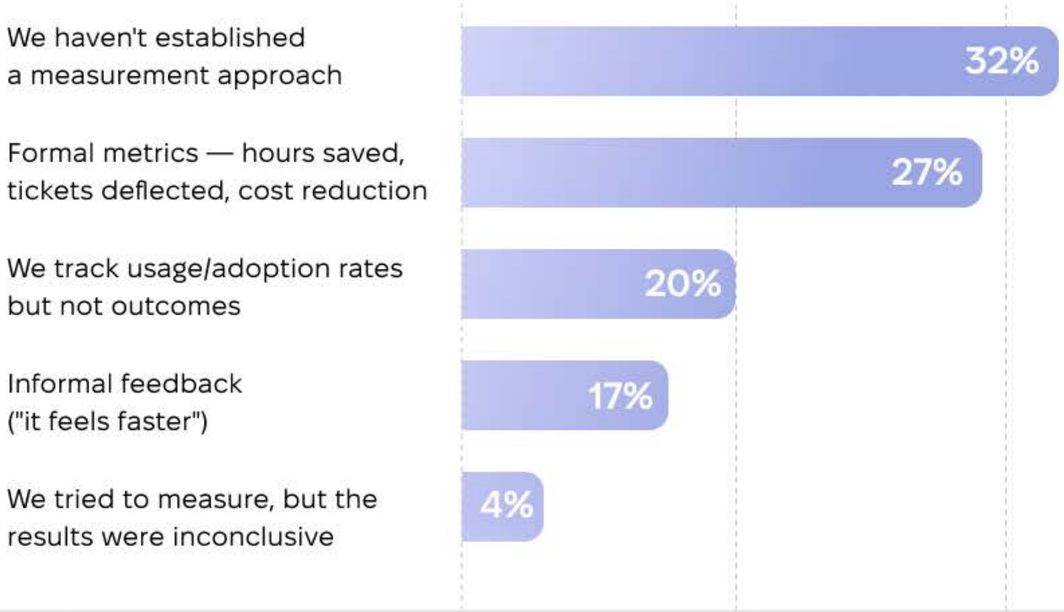

The measurement gap by company size:

Medium companies are the most exposed:

56% have not established any measurement approach

40% of mid-market

18% of enterprise

Enterprise organizations lead in formal metrics at 35% but that still leaves nearly two-thirds without outcome-based measurement.

How Many AI Initiatives Prove Their Value?

17%

of organizations say the majority of their AI initiatives deliver measurable outcomes.

27%

have no reliable way to measure AI results at all.

Key Finding: Measurement Reveals Failure, Not Prevents It

Organizations with formal ROI metrics are are more often reporting stopping or scaling back initiatives (37%) than those with no measurement approach (16%).

Measurement does not prevent failure — it reveals it. Without it, underperforming initiatives tend to linger without clear direction.

Where AI is Quietly Delivering Value

We asked every respondent to name their single most impactful AI use case. The answers were consistent: the use cases that work are specific, narrow, and task-level.

Nobody described enterprise-wide AI transformation delivering results. They described AI helping a specific team do one thing faster, with a human still in the loop.

Examples by Theme:

* subjective evaluation of our respondents

Document Processing & Knowledge Extraction

- Reviewing thousands of technical documents → 30–40% time savings
- Transferring knowledge from retiring employees to new hires

Engineering & Code Review

- Coding assistants speed up code review — developers still verify output
- Reducing engineering time for internal tooling and product development

Customer Support & Operations

- Smart search and word recognition reduce tickets by ~33%
- Help desk bots improve assignment and resolution speed

Compliance & Back-Office Automation

- Automating onboarding and legal reviews → ~70% efficiency gains
- AI tools flag missing or inconsistent application data for human verification

Highlighted Quotes

"The generated code is often poor, but it does a good job reviewing human-written code."

Senior Manager, IT Engineering

“Results are sent to humans to verify the work, especially the recommended denials.”

IT Director, Mid-Market

“AI has helped the engineering team save 30–40% of their time spent on research.”

Director IT, Enterprise

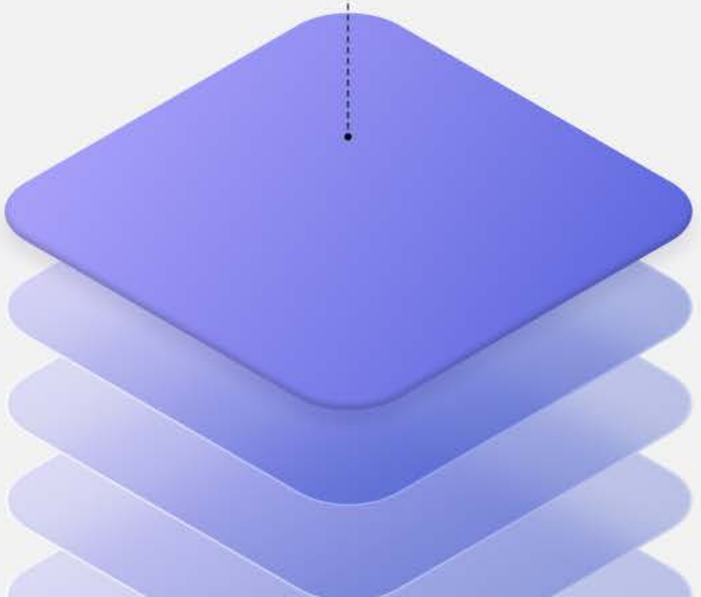

Where AI Initiatives Are Falling Short

Despite early momentum, many AI initiatives are struggling to sustain results.

More than a quarter of organizations have paused, scaled back, or fully abandoned at least one AI initiative in the past 12 months. Another 23% report initiatives that are underperforming and at risk.

In the past 12 months, has your organization paused, scaled back, or abandoned an AI initiative that was already underway?

- All initiatives are progressing as planned
- Haven't had formal initiatives long enough
- Initiatives not delivering, may be at risk
- Fully stopped at least one initiative
- Significantly scaled back scope or budget

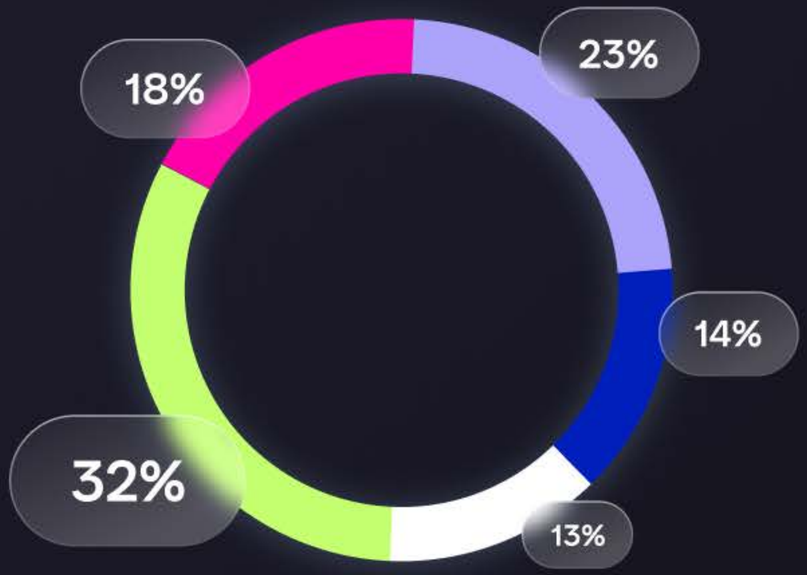

The Data Problem AI Inherited

Most AI accuracy problems aren't AI problems. They're documentation problems that existed long before anyone deployed a model.

90% of organizations deployed AI without extensively auditing the internal documentation that those tools would access.

Before deploying AI tools, did your organization audit or clean up the internal documentation AI tools would access?

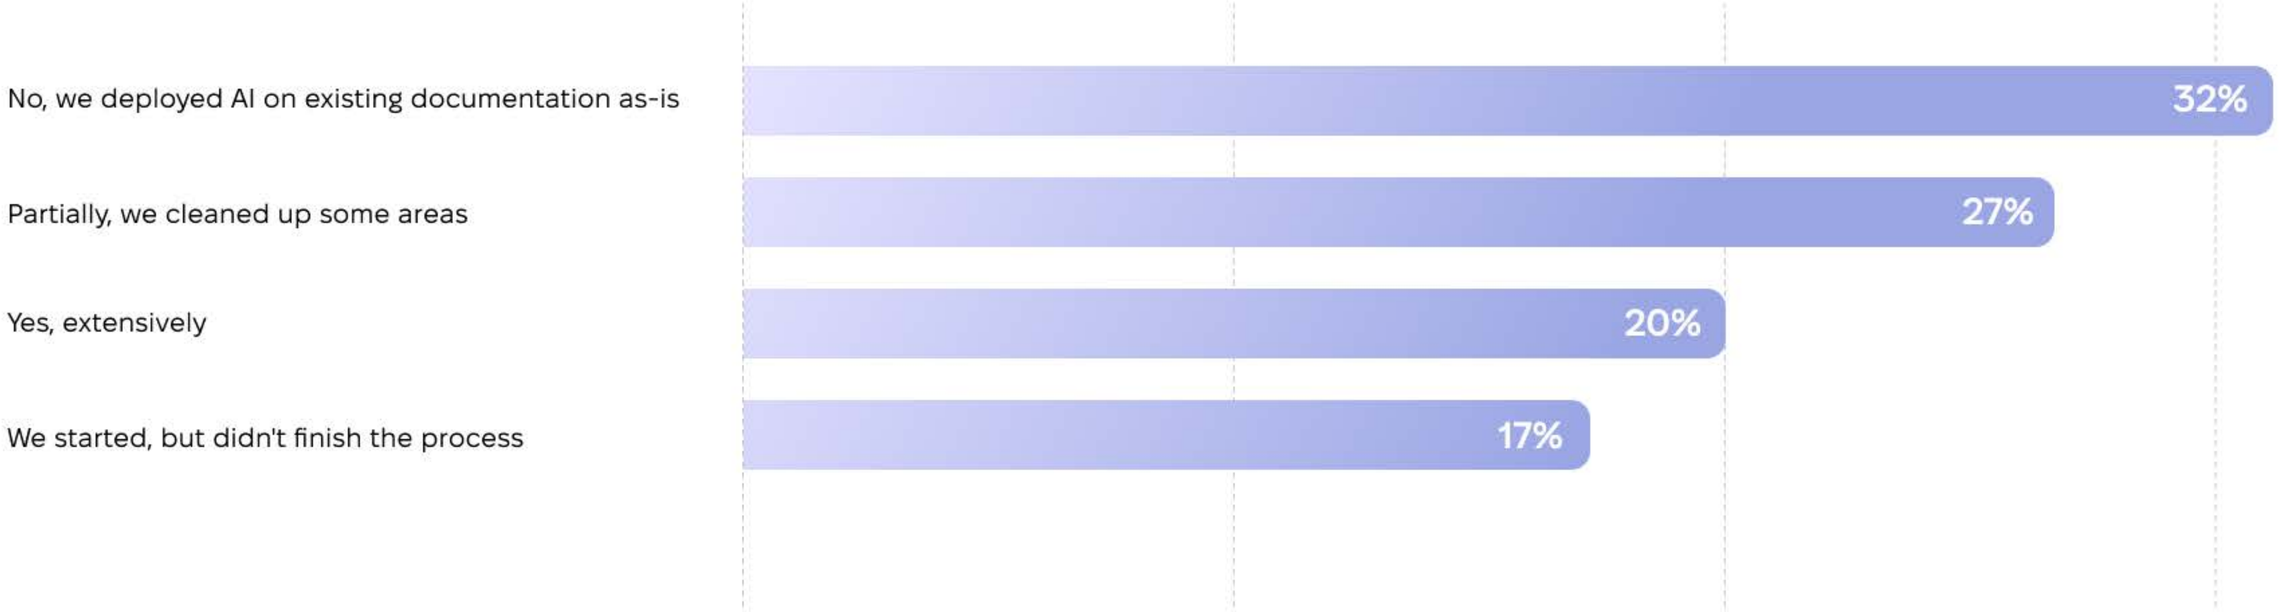

The audit gap by company size:

Medium companies are the most exposed.

62%

deployed AI on existing documentation as-is

4%

audited extensively

Enterprise organizations fare somewhat better.

37%

still deployed without any cleanup

14%

audited extensively

7%

who started the cleanup process and didn't finish it tell their own story: this work is important enough to begin, but unglamorous enough to abandon.

What Your AI Is Actually Reading

When asked to estimate what percentage of their internal knowledge base is current and accurate, respondents painted a concerning picture.

What percentage of your internal knowledge base would you estimate is current and accurate?

Estimated accuracy

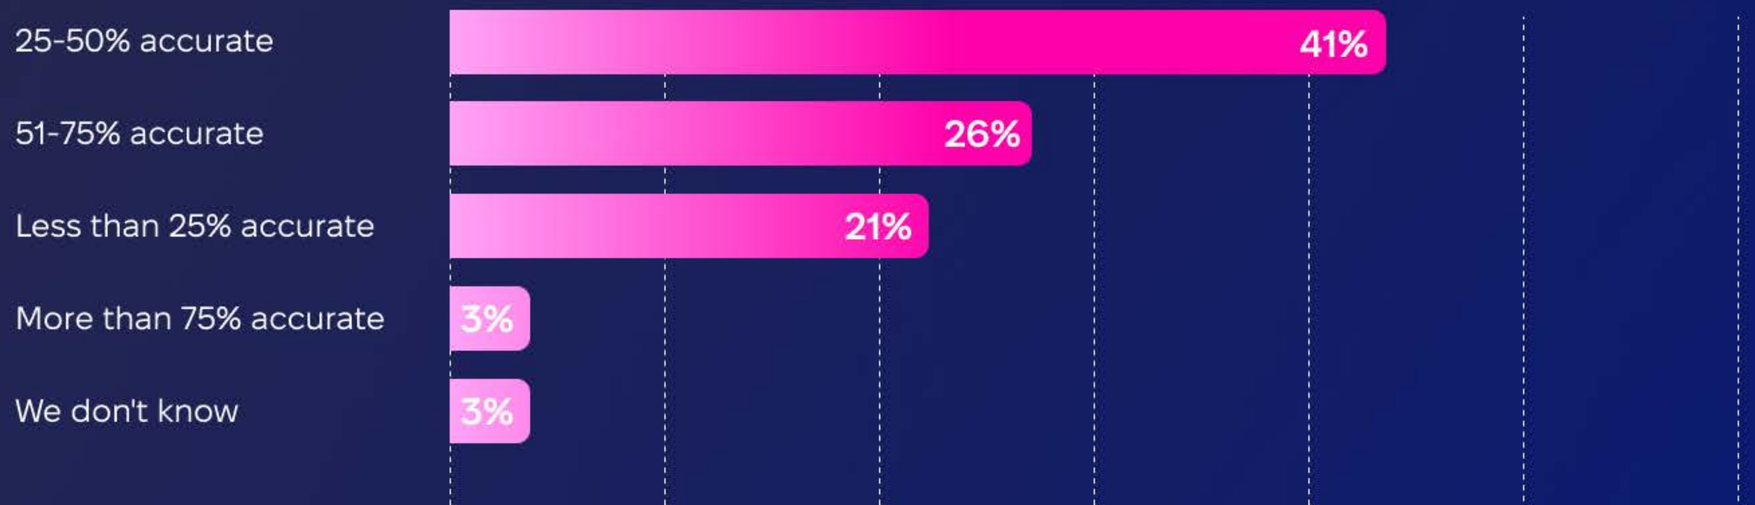

32%

estimate their knowledge base is less than 50% accurate

10%

audited their documentation extensively before deploying AI

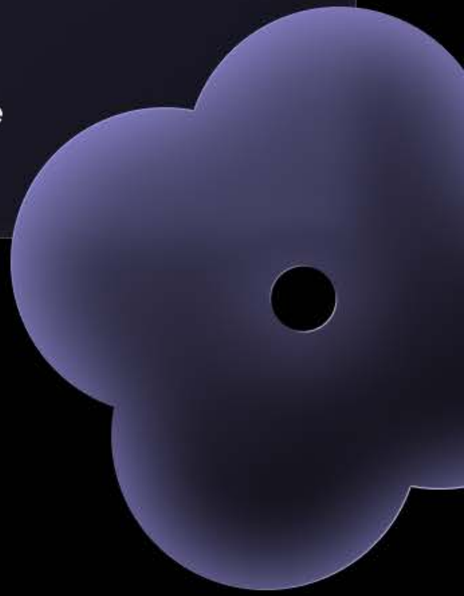

Nearly 4 in 10 organizations are either feeding AI tools documentation that is more wrong than right, or have no visibility into what they are feeding it. Among medium companies.

15%

selected "we don't know," the highest of any segment, compared to just 2% at enterprise.

How AI Adoption is Outpacing Control

The way most organizations adopt AI is the same way most organizations lose control of it.

43% are running a mixed model where IT governs some tools while teams deploy others independently, often without oversight.

Before deploying AI tools, did your organization audit or clean up the internal documentation AI tools would access?

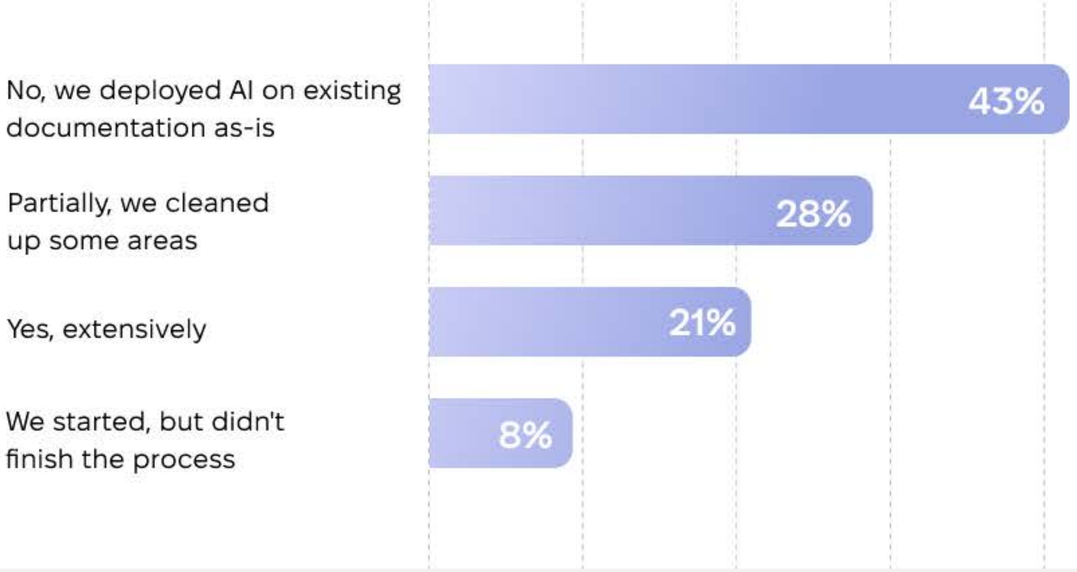

The Mid Market Adoption Paradox:

56%

of mid-market organizations are running a mixed adoption model, the highest of any segment. Not a single mid-market respondent (0%) reported a fully decentralized model.

This suggests IT teams are aware of the risk but lack the resources to fully centralize. The result: a governed layer that IT can see, and a shadow layer that IT knows exists but cannot fully manage.

Do Organizations Actually Trust AI?

98% of respondents verify AI outputs before acting on them.

The industry push towards autonomous agents is running ahead of organizational trust – for good reason

How much do you trust AI-generated outputs in your organization today?

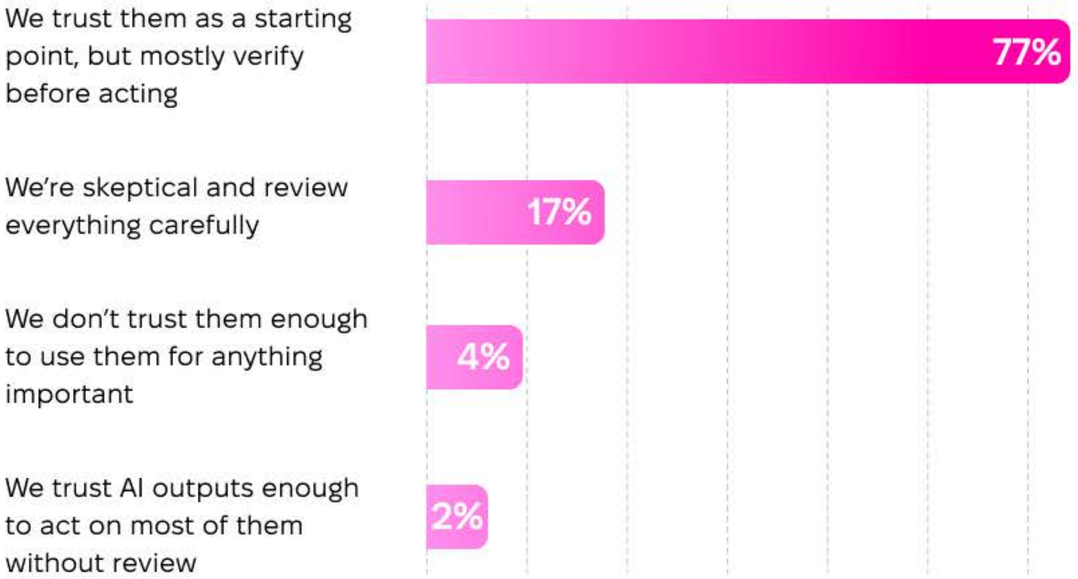

Not a single mid-market respondent trusts AI enough to act without human review 0% vs. 2% enterprise

77%

who verify as a starting point represent the current operating norm: AI as a capable first draft that humans refine, not an autonomous agent that humans supervise.

Key Finding: The Gap Between the Industry Narrative and Operational Reality

While the industry pushes toward autonomous agents, only 2% of organizations trust AI enough to act without review.

Most are still operating in a human-in-the-loop model, and for good reason.

The Security Blind Spots Emerging

21%

AI adoption is introducing new risks, often without full visibility.

21% of organizations have experienced a security incident or near-miss, or have no visibility into how AI tools are being used.

Has your organization experienced a security incident related to AI tool usage?

- No incidents, but we've identified unaddressed risks
- No, no incidents reported or risks identified
- Not sure because we don't have visibility into how AI tools are being used
- Almost, a near miss that was caught before damage
- Yes, a confirmed incident

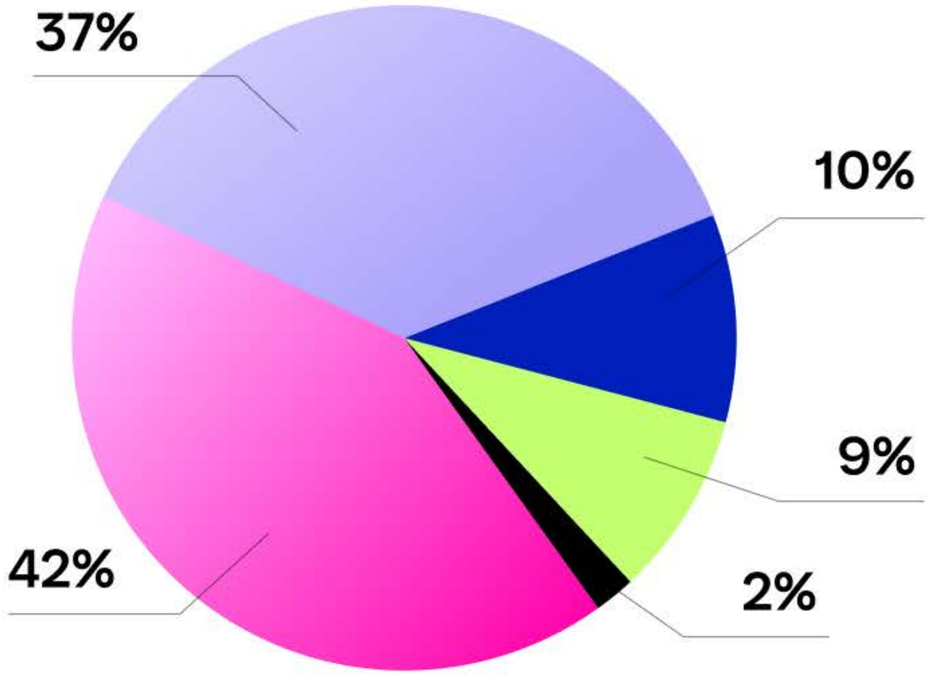

Security Exposure by Company Size:

Mid-market organizations show the highest combined rate of incidents, near-misses, and blind spots.

12%

reported near-misses or confirmed incidents

12%

have no visibility into AI usage

Medium companies have the highest "no visibility" rate

15%

suggesting that smaller IT teams simply cannot monitor tool usage at scale.

What Separates the Teams Getting it Right

There is a gap forming between the two types of organizations, and it's widening quickly.

On one side:

teams that started with governance, cleaned up their documentation, set expectations with leadership, and defined what success looks like before deploying anything. These are the minority. But their initiatives are on track, their failures are caught early, and their AI investments are producing measurable results.

On other side:

teams that deployed first and are now working backward, trying to bolt on governance, clean up data, and justify ROI after the tools are already live. These teams are not failing because they chose the wrong technology. They are failing because the organizational scaffolding was never built.

The uncomfortable truth

in this data is that most organizations fall into the second group, and the distance between the two is growing. The teams with alignment, measurement, and clean data are compounding their gains. The teams without them are compounding their problems: shadow tools they cannot see, stale documentation feeding bad answers, and leadership expectations that drift further from reality each quarter.

This is not a call to slow down.

The pressure to adopt AI is real, and the organizations that figure it out will have a significant advantage. But the data is clear that figuring it out starts with the work nobody wants to do: auditing 20 years of Confluence pages, defining what a successful outcome actually looks like, and having an honest conversation with leadership about what is possible with current resources.

Before the next AI initiative, ask:

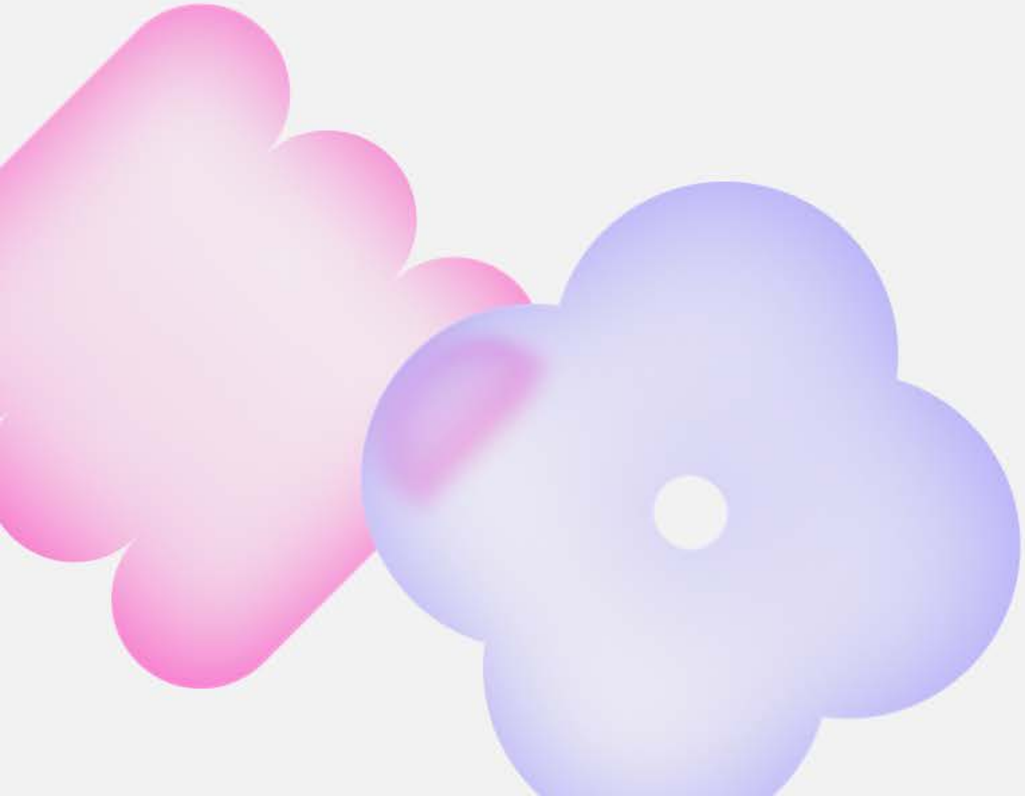

Do we know what our AI is reading?

When was our documentation last reviewed, and can we trust it?

How will we measure success?

Do we have clear, outcome-based metrics in place from day one?

Are expectations aligned with reality?

Does leadership understand what is achievable with current resources?

Do we have full visibility into AI usage?

Can we see which tools are being used and where?

Are we starting narrow enough?

Are we focusing on specific, high-impact use cases rather than broad transformation?

About Leebry

Leebry is a Work AI platform designed to help organizations make better use of the knowledge and tools they already have. By connecting systems like Confluence, Slack, Jira, Okta, and HRIS, it provides secure, citation-backed answers and enables everyday workflows through a single interface.

Leebry is built by MacPaw, known for products like CleanMyMac, Setapp, and ClearVPN, as part of its move into enterprise AI.

Available Q2 2026

Learn more at leebry.com ↗

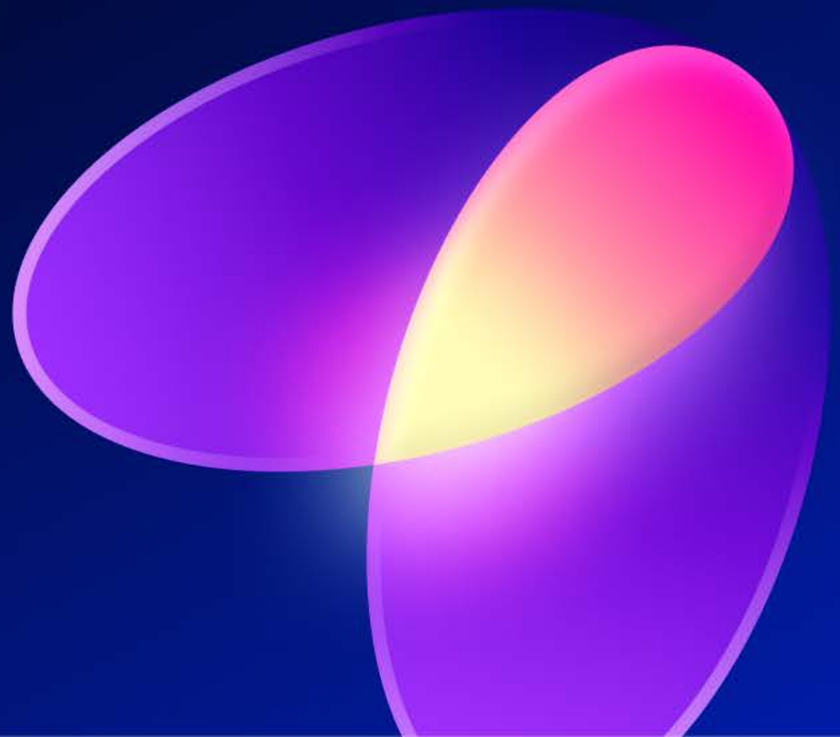

Leebry
by MacPaw