

AI agents in enterprise: *real-world use cases* & ROI

From pilot to production — autonomous AI systems are reshaping how businesses operate, cutting costs, and generating measurable returns across every industry.

\$7.6B
AI agents market size in 2025

\$183B
Projected by 2033 · CAGR ~50%

79%
of enterprises have adopted AI agents in some form

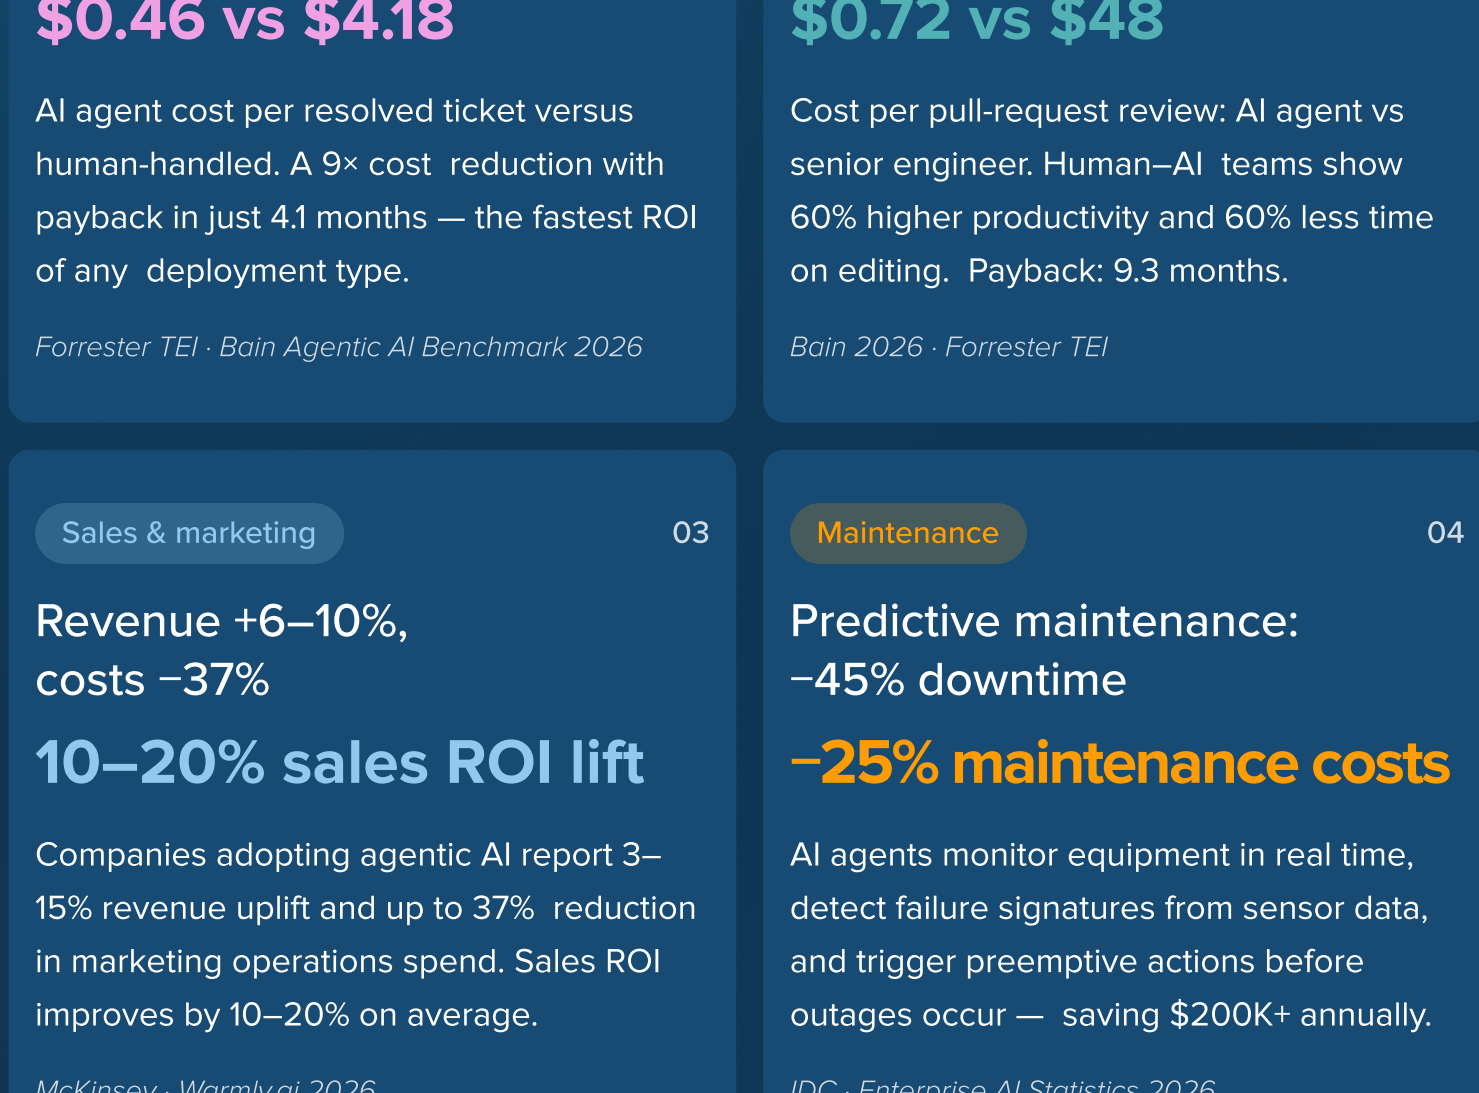
40%
of enterprise apps will embed AI agents by end of 2026

Globe Newswire · Gartner · Landbase 2026

01 What is an AI agent?

| | | |
|--|---|---|
| <p>TYPE 01</p> <p>Task agent</p> <p>Executes a single task autonomously — answering a support ticket, generating a report, or summarising documents without human intervention.</p> <p>Market share: 59% of deployments</p> | <p>TYPE 02</p> <p>Workflow agent</p> <p>Orchestrates multi-step processes across systems — CRM → email → Slack → ERP — eliminating manual hand-offs between tools.</p> <p>Growing fastest in enterprise</p> | <p>TYPE 03</p> <p>Multi-agent system</p> <p>Teams of specialised agents collaborate under central coordination to solve complex, end-to-end business processes autonomously.</p> <p>Represents 66.4% of the market</p> |
|--|---|---|

02 Adoption by industry



93%
of IT leaders plan to deploy autonomous agents within 2 years — *MuleSoft/Deloitte*

50%
of GenAI enterprises will deploy AI agents by 2027, up from 25% today — *Deloitte*

03 Real-world use cases

| | |
|---|---|
| <p>Customer service 01</p> <p>ServiceNow: –52% case handling time \$0.46 vs \$4.18</p> <p>AI agent cost per resolved ticket versus human-handled. A 9× cost reduction with payback in just 4.1 months — the fastest ROI of any deployment type.</p> <p><i>Forrester TEI · Bain Agentic AI Benchmark 2026</i></p> | <p>Software dev 02</p> <p>Code review: 66× cost reduction per PR \$0.72 vs \$48</p> <p>Cost per pull-request review: AI agent vs senior engineer. Human–AI teams show 60% higher productivity and 60% less time on editing. Payback: 9.3 months.</p> <p><i>Bain 2026 · Forrester TEI</i></p> |
|---|---|

| | |
|---|--|
| <p>Sales & marketing 03</p> <p>Revenue +6–10%, costs –37% 10–20% sales ROI lift</p> <p>Companies adopting agentic AI report 3–15% revenue uplift and up to 37% reduction in marketing operations spend. Sales ROI improves by 10–20% on average.</p> <p><i>McKinsey · Warmly.ai 2026</i></p> | <p>Maintenance 04</p> <p>Predictive maintenance: –45% downtime –25% maintenance costs</p> <p>AI agents monitor equipment in real time, detect failure signatures from sensor data, and trigger preemptive actions before outages occur — saving \$200K+ annually.</p> <p><i>IDC · Enterprise AI Statistics 2026</i></p> |
|---|--|

| | |
|--|---|
| <p>Finance 05</p> <p>BlackRock & S&P: agent teams in production Days → minutes</p> <p>Coordinated AI agent teams handle portfolio management and risk analytics. Workflows that once took days now collapse to minutes — at a scale humans can't match manually.</p> <p><i>Snowflake Data Trends 2026</i></p> | <p>Maintenance 06</p> <p>Clinomic Mona: real-time ICU support 56% faster AI adoption</p> <p>AI agents consolidate, analyse and visualise patient data for ICU teams in real time. Medical AI publication volume has more than doubled in two years — fastest-growing sector.</p> <p><i>NVIDIA State of AI 2026 · Stanford AI Index</i></p> |
|--|---|

04 Companies leading the way

| |
|---|
| <p>Klarna. Fintech · Sweden</p> <p>Customer service AI agent handling 2.3M conversations/month</p> <p>Klarna deployed an OpenAI-powered agent as the primary first-contact layer across all customer service channels. The agent can process refunds, resolve disputes, check order status, and answer policy questions — all via live API connections to internal systems, not just FAQ responses. It handles roughly two-thirds of all support chats without any human involvement.</p> <p>Resolution time: 11 min → 2 min 853 FTE equivalent workload</p> <p>CSAT on par with human agents 25% fewer repeat contacts</p> <p>\$60M saved annually. AI does the work of 853 full-time agents.</p> <p><i>Lesson learned: Klarna later reintroduced human agents for or complex disputes, settling on a hybrid model. Full automation works for high-volume, rule-based queries — nuanced issues still need humans.</i></p> |
|---|

| |
|--|
| <p>JPMorganChase Banking · USA</p> <p>LLM Suite: from research tool to enterprise agent orchestration layer</p> <p>JPMorgan's proprietary LLM Suite started as an internal research assistant but has evolved into a multi-agent orchestration platform. AI agents generate M&A presentations in 30 seconds (vs hours for junior analysts), automate trade settlement workflows, detect fraud in real time, and increasingly perform "computer-use" actions — navigating internal software the way a human employee would. \$18B annual technology budget is channelled into this infrastructure.</p> <p>M&A memos: hours → 30 seconds</p> <p>\$18B annual tech investment Real-time fraud detection</p> <p>450+ active AI agent use cases running in production daily.</p> <p><i>JPMorgan's approach: treat agents like licensed professionals — every action requires audit trails, permission enforcement, and explainability logs before it can scale across the enterprise.</i></p> |
|--|

| |
|---|
| <p>salesforce Enterprise SaaS · USA</p> <p>Agentforce: 380,000 customer interactions handled autonomously</p> <p>Salesforce deployed its own Agentforce platform internally for customer support. The system handled over 380,000 interactions — answering product questions, processing requests, and routing issues — with an 84% autonomous resolution rate. A Fortune 500 client using Agentforce for internal reporting cut report generation time from 15 days to 35 minutes, reducing cost per report from \$2,200 to just \$9.</p> <p>R380K interactions handled Report time: 15 days → 35 min</p> <p>Cost per report: \$2,200 → \$9</p> <p>84% of cases resolved autonomously. Only 2% escalated to humans.</p> <p><i>The reporting automation case is among the most striking cost-reduction examples in enterprise AI — a 99.6% reduction in per-report cost with no loss of accuracy.</i></p> |
|---|

| |
|---|
| <p>Walmart Retail · USA</p> <p>Autonomous inventory and demand planning across 4,700 stores</p> <p>Walmart's supply chain AI agent ingests historical and real-time sales data from all stores and fulfillment centres, then makes autonomous replenishment decisions without human approval loops. The agent continuously processes demand signals, adjusts forecasts, and triggers stock transfers — at a scale impossible for human planners. A comparable North American retailer using similar agents reduced quarterly inventory losses from \$5.4M to \$1.6M.</p> <p>Zero per-decision human sign-off Inventory loss: \$5.4M → \$1.6M</p> <p>Real-time across 4,700 locations</p> <p>4,700 stores connected to one autonomous supply chain forecasting agent.</p> <p><i>The reporting automation case is among the most striking cost-reduction examples in enterprise AI — a 99.6% reduction in per-report cost with no loss of accuracy.</i></p> |
|---|


| |
|---|
| <p>BANK OF AMERICA Banking · USA</p> <p>Erica virtual assistant: proactive financial guidance at scale</p> <p>Bank of America's AI assistant Erica has surpassed 3 billion client interactions globally, handling tens of millions per month. Unlike early chatbots that deflected queries, Erica proactively surfaces financial insights, alerts users to unusual spending, and helps navigate products — reducing call centre volume while increasing engagement. It is deeply integrated into the core mobile banking app used by over 47 million active clients.</p> <p>3B+ total interactions Reduced call centre volume</p> <p>Proactive (not just reactive) AI</p> <p>3B+ client interactions handled by Erica since launch.</p> <p><i>Erica's success stems from proactive design — surfacing insights without being asked — rather than just answering queries. This shifts AI from a cost centre to a customer engagement driver.</i></p> |
|---|

| |
|--|
| <p>AtlantiCare Healthcare · USA</p> <p>Clinical AI assistant cuts documentation time by 42% across ICU providers</p> <p>AtlantiCare deployed an agentic AI clinical assistant that listens to patient encounters and automatically generates structured clinical notes, freeing physicians from the documentation burden that accounts for 35–40% of a doctor's working day. The system achieved 80% adoption among the 50 test providers in the pilot group — an unusually high rate for clinical technology — and has since expanded across the organisation.</p> <p>Documentation time –42% 80% clinician adoption rate</p> <p>66 min/day returned per doctor</p> <p>66 min freed per clinician per day through AI documentation automation.</p> <p><i>The reporting automation case is among the most striking cost-reduction examples in enterprise AI — a 99.6% reduction in per-report cost with no loss of accuracy.</i></p> |
|--|

05 ROI benchmarks

| | | |
|--|---|--|
| <p>171% Average ROI from agentic AI. US enterprises: 192%. That is 3× the return of traditional automation.</p> | <p>6.4 hrs Saved per knowledge worker per week. <i>McKinsey Global AI Survey 2026.</i></p> | <p>\$1.49 Returned for every \$1 invested in GenAI. <i>Snowflake / Omdia 2026.</i></p> |
| <p>5.8× ROI within 14 months of production deployment — <i>McKinsey.</i></p> | <p>62% of companies expect 100%+ ROI from their AI agent deployments.</p> | <p>38 days Median time to first value — vendor agents vs 94 days custom builds. <i>Bain 2026.</i></p> |

MEDIAN PAYBACK PERIOD BY DEPARTMENT



06 What holds companies back

- Only 41% hit ROI in year one**

19% of projects never reach payback — primarily due to evaluation drift and governance gaps, not agent capability. The 90% pilot-to-production accuracy drop is the #1 cited cause. (*Gartner Agentic AI Pulse 2026*)
- 40% of projects at risk of cancellation by 2027**

Only 21% of organisations have a mature governance model for autonomous AI agents. Ungoverned deployments compound quickly. Projects cancelled in 2027 are being built without governance today. (*Gartner*)
- 52% cite data quality as their #1 blocker**

IDC: organisations that fail to establish AI-ready data foundations will lose an estimated 15% of productivity by 2027. Bad data doesn't produce a bad report anymore — it produces a bad action, at machine speed.
- 79% adopted — but only 11% run agents in production**

The largest deployment backlog in enterprise technology history. The gap between organisations that have moved to production and those still in pilots is widening — and the window to catch up without competitive disadvantage is narrowing.

07 Key takeaways for enterprises

| | |
|---|---|
| <p>Start with proven ROI zones</p> <p>Customer service, marketing ops, and code review offer the fastest payback and lowest implementation risk. Build internal credibility before expanding to complex multi-agent systems.</p> | <p>Fix your data before scaling</p> <p>Over half of all failures trace back to poor data quality. Enterprises that invest in AI-ready data foundations first see substantially better outcomes across every deployment type.</p> |
|---|---|

| | |
|--|--|
| <p>Governance is not optional</p> <p>Real-time monitoring, audit trails, human-in-the-loop controls, and kill switches dramatically reduce programme cancellation risk. Build governance in from day one — not after the first failure. Only 21% of enterprises have a mature governance model today; that gap is the single biggest predictor of project cancellation by 2027.</p> | <p>Match the build approach to the complexity</p> <p>For high-volume, well-defined tasks — ticket routing, FAQ resolution, basic reporting — vendor tools get you to value in 38 days. For complex, domain-specific workflows, build custom from day one: retrofitting a vendor solution to your data model and compliance requirements almost always takes longer than a purpose-built agent designed around your architecture from the start.</p> |
|--|--|

Partner with a vendor who has done it before

The gap between a successful AI agent deployment and a stalled pilot almost always comes down to one thing: implementation experience. The enterprises hitting 171% ROI didn't figure it out alone — they worked with partners who had already solved the data, governance, and integration challenges across multiple deployments. Choose a team that brings battle-tested playbooks, not just technical capability.

● InData Labs — AI & data engineering since 2014 [See our AI agent projects →](#)

INDATA LABS · AI CONSULTING

Ready to deploy AI agents in your organisation? [Explore our services →](#)

We design, build and scale custom agentic AI solutions — from strategy to production.