

Contents

Foreword	3
Research insights	6
Part I: The current landscape	
How organizations are deploying AI agents today	7
Part II: Going deeper	
Expanding use cases and measuring ROI	12
Part III: The path forward	
Driving AI agent adoption in 2026	18
AI agents in production	23
Outlooks and perspectives	38
Additional resources	45

Foreword

Over the past several months, AI agents have moved from experimental technology to infrastructure that enterprises use in production. Unlike traditional software that waits for human input, agents reason through problems, make decisions, and take action autonomously—handling everything from multi-step coding workflows to cross-functional business processes.

This shift toward automated workflows and multi-step agentic systems fundamentally changes what organizations require from AI: models that are secure when handling proprietary data, compliant with industry regulations, and robust against adversarial attacks like jailbreaks.

In partnership with research firm [Material](#), we surveyed over 500 technical leaders in the United States across company sizes and industries to understand how organizations are using agents today and where they see opportunity in 2026. What emerged is a clear picture of technology in transition: from task automation to strategic impact, from single-function pilots to cross-functional deployment, and from incremental efficiency to fundamental shifts in how work gets done.

The data shows this shift in concrete terms. According to our research, more than half of organizations (57%) now deploy agents for multi-stage workflows, including 16% that have progressed to cross-functional processes spanning multiple teams. In 2026, 81% plan to tackle more complex use cases—39% developing agents for multi-step processes and 29% deploying them for cross-functional projects.

Given the growth of agentic coding over the past 12 months, it's unsurprising that nearly 90% of organizations surveyed use AI to assist with coding today.

Organizations report AI agents free up more time across the entire development lifecycle—from planning and ideation (58%) to code generation, documentation, testing, and review (all at 59%).

The impact also extends well beyond software development. Beyond engineering, the highest-impact use cases include data analysis and report generation (60%) and internal process automation (48%), with 56% planning to implement agents for research and reporting over the next year. And 80% report these investments are already delivering measurable economic returns—not projected value or pilot results, but actual ROI.

Eight in 10 organizations believe AI agents have already delivered measurable ROI, with another 1 in 10 saying they expect them to deliver more economic impact in the future. The question facing leaders in 2026 isn't whether to adopt AI agents but how to scale them strategically while addressing integration challenges (46%), data quality requirements (42%), and change management needs (39%).

Read on to learn more about how today's leaders are building AI agents in the enterprise.

Survey methodology

In partnership with Material, Anthropic surveyed over 500 technical leaders across company sizes and industries in late 2025 to understand current AI agent adoption patterns and future plans. Respondents included engineering leaders, IT executives, and technical decision-makers from organizations ranging from startups to large enterprises across multiple sectors.



Research insights

Research insights

Part I: The current landscape

How organizations are deploying AI agents today

More than half of businesses are deploying AI agents with multi-step workflows


Trend summary

Organizations are deploying AI agents for work that goes well beyond chat interfaces and single-step automation. More than half (57%) now use agents to handle multi-stage workflows, while 16% have progressed to cross-functional or end-to-end processes spanning multiple teams or business functions.

Why this matters

The shift from task automation to process orchestration represents a fundamentally different use case—and a different value proposition. Organizations that master multi-stage and cross-functional agent deployments can unlock advantages in speed, consistency, and scale that simple automation can't. This is where AI moves from incremental efficiency gains to enabling new ways of working.

What type of AI agents are you deploying?




Nearly all organizations are adopting coding agents

Trend summary

More than 9 in 10 organizations now use AI to assist with coding. The vast majority (86%) have moved beyond experimentation and are deploying AI coding agents for production code, with enterprises leading adoption at 91% compared to 83% for small and mid-sized businesses. And 42% of organizations trust these agents to lead development work with human oversight, signaling a meaningful shift in how engineering teams are structured and how code gets written.

Why this matters

AI coding agents have moved from experimental to mainstream, with the majority of organizations already deploying them in production environments. Organizations that embrace these tools strategically are accelerating delivery timelines, optimizing engineering resources, and freeing developers to focus on higher-value architectural and problem-solving work. The split between organizations that trust agents to lead versus assist reveals significant upside for those who invest in building expertise and establishing best practices early.



AI coding agents boost developer productivity


Trend summary

AI agents are increasing productivity across the entire development lifecycle, not just code generation. Organizations report time gains in four key areas at nearly identical rates: code generation (59%), research and documentation (59%), code review and testing (59%), and planning and ideation (58%).

Why this matters

The impact spans every phase of software development, which means teams can improve both engineering velocity and code quality simultaneously. Organizations that integrate AI agents across the full development process can compound these gains, turning what might be a 10-15% improvement in coding speed into meaningful acceleration of entire project timelines. The nearly even distribution across activities also suggests teams are finding value wherever they apply these tools, making this less about picking the "right" use case and more about systematic adoption.

Where do AI agents open up more time for engineers across the software development lifecycle?




Leaders favor hybrid approaches to agentic development over building from scratch

Trend summary

Most organizations (47%) take a hybrid approach to AI agents, combining off-the-shelf solutions with custom-built components. About one in five (21%) rely entirely on pre-built agents, while a similar share (20%) build their own using APIs, open-source models, or developer toolkits that require coding expertise.

Why this matters

The hybrid model dominance suggests no single approach delivers everything organizations need. Off-the-shelf agents get teams running quickly but often lack the customization required for specific workflows or proprietary systems. Fully custom builds offer control and differentiation but require significant engineering investment. Most organizations find value in the middle: using pre-built agents where they work well and investing development resources only where customization creates meaningful advantage.



Research insights

Part II: Going deeper

Expanding use cases and measuring ROI

Agentic use cases are expanding beyond coding


Trend summary

Organizations expect AI agents to expand well beyond engineering and IT functions over the next 12 months. Research and reporting leads adoption plans at 56%—particularly among mid-market and enterprise organizations—followed by supply chain optimization, product development, and financial planning. The breadth of planned use cases signals a shift toward treating AI agents as enterprise-wide infrastructure rather than department-specific tools.

Why this matters

Research and reporting work spans every function and level of an organization, making it a high-leverage starting point that builds institutional comfort with AI agents before deploying them in more sensitive or complex workflows. Organizations that successfully implement agents for research and analysis can establish governance frameworks, build internal expertise, and demonstrate ROI in ways that accelerate adoption for higher-stakes use cases like financial planning or supply chain decisions. The cross-functional nature of these early deployments means capabilities compound across the business, not just within isolated teams.

What types of agents does your organization intend to build in the next 12 months?



In addition to coding, data analysis and process automation are the enterprise’s most impactful agentic use cases

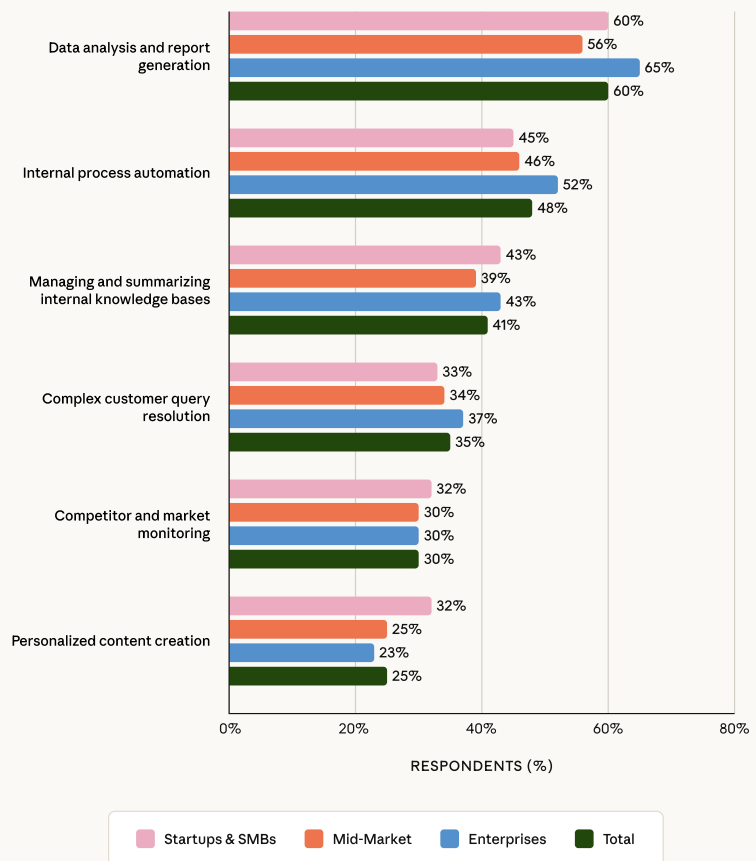
Trend summary

Beyond coding, the highest-impact AI agent use cases are data analysis and report generation (60% say this is one of the most impactful tasks) and internal process automation (48%). Enterprises are particularly bullish on data analysis and reporting, with 65% citing these as high-impact applications

Why this matters

Data analysis and reporting work touches every part of an organization—finance needs monthly reports, sales needs pipeline analysis, operations needs supply chain visibility. The enterprise enthusiasm is telling, as larger organizations typically have more data, more complex reporting requirements, and more people spending time on analysis work that agents can accelerate or automate entirely. Internal process automation delivers a different kind of value, reducing friction in repetitive workflows that slow teams down but don't require deep expertise. Organizations should prioritize use cases where agents can either amplify expert judgment (data analysis) or eliminate low-value work (process automation), rather than simply digitizing existing manual processes.

Which of the following agentic tasks have been the most impactful for your organization?



Efficiency gains are the primary unlock from AI agents


Trend summary

Organizations expect AI agents to deliver efficiency gains over the next 12 months, with 44% anticipating faster task completion. Enterprises also anticipate an additional benefit beyond velocity: measurable cost savings from their agent deployments.

Why this matters

The split between efficiency gains and cost savings reveals two distinct paths for AI agents today, and both create room for what comes next. Speed improvements help organizations do more with existing resources while cost savings, which enterprises are particularly positioned to capture at scale, come from reducing manual effort and avoiding expensive errors. As organizations mature their agent deployments, these gains unlock entirely new categories of work: comprehensive competitive analysis, continuous documentation, proactive customer outreach—efforts that weren't economically viable before. The organizations capturing the most value in 2026 will be pursuing opportunities that only exist because of compounding efficiency gains.

What measurable outcomes does your organization expect from AI agents over the next 12 months?



Leaders expect AI agents to drive ROI across the company


Trend summary

In 2026, software development (57%) and customer service (55%) are expected to see the greatest near-term impact from AI agents, with marketing and sales (46%) and supply chain, logistics, and operations (44%) close behind.

Why this matters

These four functions share key characteristics that make them ideal proving grounds for AI agents: they involve high-volume repetitive work, require fast iteration cycles, and have clear performance metrics that make ROI measurable. The proximity in expected impact across these functions—ranging from 44% to 57%—suggests we’re seeing multiple viable entry points rather than one dominant use case.

In what functional areas do you expect AI agents to have the biggest impact at your organization over the next 12 months?



80% of leaders say AI agents are delivering financial value today

Trend summary


The majority of organizations (80%) report that their AI agent investments are already delivering measurable economic impact today, and confidence is even higher looking forward—88% expect continued or increased returns. This isn't speculative ROI; most organizations are seeing concrete business value from their deployments.

Why this matters


Organizations have moved past the proof-of-concept phase and into measurable returns, shifting the conversation from "should we invest?" to "how do we scale what's working?" These findings suggest that returns compound as organizations deploy agents across more use cases, refine their implementations, and build institutional knowledge. Early movers are building the expertise and infrastructure that will let them capture disproportionate value as the technology continues to mature.

To what extent to you agree or disagree with the following statements?

Our investment in AI agents **has already delivered** measurable financial impact for our company



Our investment in AI agents **is expected to deliver** measurable financial impact for our company in the future



Research insights

Part III: The path forward

Driving AI agent adoption in 2026

Leaders are optimistic about the business impact of AI agents

Trend summary

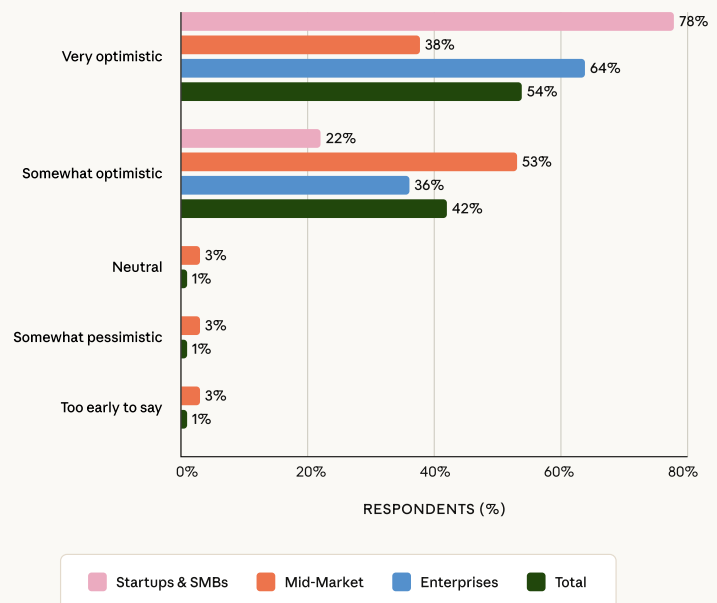
Organizations across all segments expect AI agents to deliver meaningful business impact in 2026, with enterprises showing especially strong confidence in the technology's potential.

Why this matters

Enterprise optimism is a significant signal because larger organizations typically move more cautiously—they have longer evaluation cycles, stricter governance requirements, and higher bars for proof of value. When enterprises express strong confidence, it suggests they're seeing results at scale, not just in pilots. Their bullishness also tends to influence the broader market: enterprise adoption drives vendor investment in security, compliance, and integration capabilities that eventually benefit organizations of all sizes.

How optimistic are you about the role AI agents will play at your organization in 2026?

Among decision makers who are director+ using only off-the-shelf AI agents



Data quality and integration are the biggest barriers to AI agent adoption


Trend summary

Integration and data quality challenges top the list of implementation barriers across organizations of all sizes. Nearly half (46%) cite integration with existing systems as a primary obstacle, while 42% point to data access and quality issues and 43% to implementation costs. Small and mid-sized businesses face a distinct challenge: they're significantly more likely to struggle with the human side of adoption, including employee resistance and training needs (51% compared to lower rates among larger enterprises).

Why this matters

These barriers are predictable and addressable, but they require different strategies depending on your organization's size and maturity. Enterprises need to prioritize technical integration and data infrastructure work upfront—treating AI deployment as a systems challenge, not just a software purchase. Organizations across segments that address both the technical and change management dimensions simultaneously will see faster time-to-value than those focused on technology alone.

Which of the following challenges has your organization encountered so far/do they anticipate encountering with deploying AI agents?



AI agents are shifting employee time away from rote tasks and toward strategic work

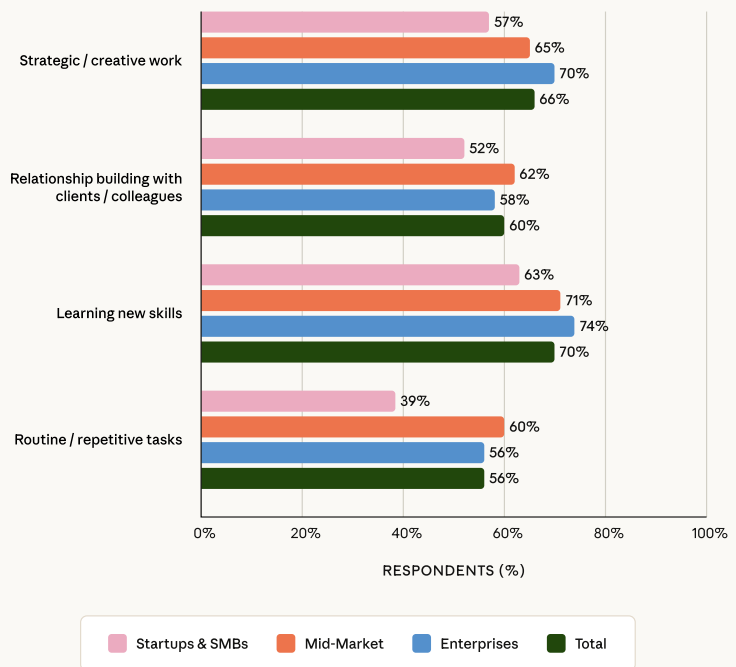
Trend summary

Agents are shifting how employees spend their time—increasing focus on strategic work (66%), relationship building (60%), and skill development (70%) rather than routine execution.

Why this matters

This addresses one of the core questions about AI adoption: whether it primarily displaces work or elevates it. The data suggests organizations are seeing the latter—agents handling execution while humans focus on judgment, relationships, and learning. That shift has implications beyond productivity metrics: teams that spend more time on strategy and skill-building become more valuable over time, not less. Organizations should design agent deployments with this goal in mind, measuring not just task completion rates but whether people are working on progressively higher-leverage problems. The companies that use agents to develop their people while improving efficiency will build sustainable advantages over those focused solely on cost reduction.

What tasks do your employees spend more time on since deploying AI agents?



8 in 10 businesses plan to implement more complex agents in 2026


Trend summary

The majority of organizations (81%) plan to move beyond simple task automation toward more complex AI projects in 2026, with enterprises leading this shift at 87% compared to 78% of SMBs. Looking at what "more complex" actually means: 39% expect to develop agents that handle multi-step processes, while 29% plan to deploy agents for cross-functional projects that span multiple teams or departments.

Why this matters

Organizations are preparing to tackle harder problems with AI—work that requires coordination across systems, functions, and decision points. The companies that identify their highest-leverage complex use cases now can build capabilities and institutional knowledge while others are still focused on basic automation. Think customer intelligence that informs sales strategy, contract lifecycle management that connects legal through procurement, or strategic planning agents that synthesize inputs from finance, operations, and product teams.

Over the next year, which best describes your organization's plans for building AI agents?



AI agents in production

AI agents have moved from pilot programs to production systems faster than most enterprises anticipated. And they're not operating as order-takers or task-managers. Leading organizations have moved agents from the periphery of operations to the center of how work gets done.

Take web development startup, **Lovable** for example. With agentic coding tools, the company now ships code 20x faster than writing it manually. **Thomson Reuters** can deliver comprehensive legal analysis pulled from 150 years of legal expertise in minutes. And cybersecurity leader **eSentire** compressed threat analysis from 5 hours to 7 minutes. And these examples just scratch the surface.

In 2025, companies have transitioned agents from efficiency plays (coding, process automation, etc.) to sources of competitive advantage (product development, research, etc.) and we anticipate adoption of long-running, multi-step agentic systems will only accelerate. The following stories show what that looks like in practice across healthcare, retail, financial services, cybersecurity, and beyond.

Healthcare and life sciences

Healthcare and life sciences organizations face a mounting tension: the need to move fast to improve patient outcomes while maintaining the highest standards for safety, privacy, and regulatory compliance. All too often, it can feel like the pressure to build agents stands in stark contrast to existing protocols. But as we're finding, the organizations succeeding with agents aren't choosing between speed and rigor—they're achieving both.

Novo Nordisk transforms clinical documentation from months to minutes

Novo Nordisk, a global pharmaceutical company and maker of Ozempic, develops innovative medicines for chronic diseases including diabetes and obesity. The company serves millions of patients worldwide with life-changing treatments that require extensive regulatory documentation before reaching patients.

The challenge: Documentation delays drug discovery

For pharmaceutical companies, the path from drug discovery to patient access is paved with paperwork. Each new treatment comes with its own mountain of required documentation: clinical study reports running hundreds of pages, highly technical device verification protocols, and guides explaining complex treatments in consumer-friendly ways.

The pinnacle of this process is the clinical study report (CSR), a document up to 300 pages long. Staff writers averaged only 2.3 CSRs per year, and the manual process remained prone to errors. With each day of delay costing up to \$15 million in potential revenue—and patients with diabetes and obesity continuing to wait—the stakes couldn't have been higher.

The solution: AI-powered documentation platform delivers regulatory-grade content

Novo Nordisk developed NovoScribe, a generative AI platform built with Claude Code on Amazon Bedrock and MongoDB Atlas, with Claude models as the frontier intelligence driving the entire system. The platform combines retrieval-augmented generation with domain expert-approved text and case-specific variables to produce accurate, compliant documentation.

Claude Code has also fundamentally changed product development at Novo Nordisk by enabling non-technical team members to prototype ideas. This ability to spin up a prototype in hours instead of days or weeks has been a strategic catalyst for their 11-person development team, allowing them to maintain a small size while dramatically expanding capabilities.

Results and impact:

- **10+ weeks to 10 minutes** for clinical study documentation production
- **95% reduction in resources** needed to create device verification protocols
- **Review cycles cut in half** as output quality improved

Doctolib cuts engineering cycle time from weeks to hours

Doctolib, Europe's leading healthcare technology company, serves 420,000 health professionals and 90 million patients across France, Germany, Italy, and the Netherlands. With a comprehensive digital healthcare ecosystem that manages everything from appointments to electronic health records, Doctolib's engineering team faced a familiar challenge: administrative development tasks consumed time that could be spent solving complex healthcare problems.

The challenge: Repeatable tasks consuming time meant for solving healthcare problems

Engineers spent significant time on repeatable administrative tasks, including writing documentation, creating tests, and addressing technical debt. The team needed to refocus on solving complex healthcare challenges rather than wrestling with routine engineering work.

The solution: Autonomous development workflow cuts infrastructure replacement from weeks to hours

After a successful pilot with 30 engineers, Doctolib rolled out Claude Code across their entire development team. Engineers self-onboard in under five minutes across any IDE. The platform team created a centralized repository of prompts, custom commands, and subagents that all developers pull during initial setup—every engineer starts with proven, reusable workflows on day one.

The team embedded Claude Code directly into their CI pipeline through headless mode, automatically opening pull requests for routine maintenance. Every code change triggers a CI job that updates technical documentation automatically, keeping technical docs current without manual intervention.

Results and impact:

- Legacy testing infrastructure replaced in **hours instead of weeks**
- **300 daily active users**, making Claude Code the most-used AI tool at the company
- Engineering ships features **40% faster** while maintaining code quality

Retail

Retail and consumer brands generate massive amounts of customer data but can't always operationalize it. AI agents help close that gap. Employees who couldn't write SQL queries yesterday are pulling complex analytics today to better understand and act on customer needs.

L'Oréal achieves 99.9% customer analytics accuracy with AI agents

L'Oréal is the world's largest cosmetics and beauty company, operating in over 150 countries with a portfolio of 37+ international brands across skincare, haircare, makeup, and fragrance.

The company is positioning itself as a tech leader in the category, combining cutting-edge research and innovation with AI, data, and digital capabilities. But its massive amounts of customer data created bottlenecks. Regional teams couldn't answer customer questions without requesting custom reports, while marketing decisions often waited on technical teams to build dashboards

The challenge: Democratizing data access without sacrificing accuracy

For a company competing across global markets, data access delays meant missed opportunities. L'Oréal needed to democratize AI capabilities across its global workforce while maintaining security and governance. A key challenge was enabling employees to access and analyze data without requiring technical expertise or custom dashboard development. Previous GenAI approaches achieved only 90% accuracy on conversational analytics—not effective enough for building user trust at scale.

The solution: Orchestrated AI agents achieve 99.9% accuracy on complex analytics queries

L'Oréal needed AI that could pull data for its employees across multiple systems simultaneously without losing accuracy. L'Oréal selected Claude for conversational analytics through rigorous testing. Claude serves as the main orchestrator of multiple specialized agents that work together to transform user questions into insights and visualizations.

When an employee asks a question, Claude coordinates with semantic API agents, data retrieval systems, and specialized agents for calculations, product master data, and geography master data. This orchestration allows Claude to determine which specialized agents to engage based on the user's question and synthesize results into clear answers.

Results and impact:

- **99.9% accuracy** on conversational analytics applications, up from 90% with previous GenAI approaches
 - **44,000 monthly users** generating 2.5 million messages monthly
 - Employees now query data directly rather than building custom dashboards
-

Technology

Digital-first companies are embedding agents directly into their products. For these organizations, agents aren't just improving operations, they're becoming the product experience itself.

Shopify transforms merchant empowerment with AI commerce assistant

Shopify powers commerce for millions of businesses worldwide, from first-time entrepreneurs to enterprise brands. The platform makes it simple for anyone to start, run, and grow a business across online stores, brick-and-mortar locations, and everything in between. But this presents a unique challenge for Shopify's customer support team, as each customer arrives with unique backgrounds and challenges. For instance, many non-technical founders launching their first store needed guidance from basic setup to sophisticated analytics queries that typically require deep knowledge of Shopify's query language (ShopifyQL).

The challenge: Scaling expert-level support across millions of merchants

Merchants needed deep technical knowledge to extract insights about their own businesses. To do so, many had to really learn how to use ShopifyQL and get pretty deep into understanding the query language to get the insights they needed. That barrier kept many entrepreneurs from, and traditional support couldn't scale to the diverse needs of its millions of global users.

The solution: Always-available AI guides merchants from setup to first sale

To address this challenge, Shopify built Sidekick, an AI-enabled commerce assistant that acts as an always-available expert for merchants. Instead of learning Shopify's complex query language, merchants simply ask questions in plain English. Sidekick translates those questions into sophisticated data analytics, pulls in relevant product information and business context, then delivers clear answers.

The system works fast enough to feel conversational. When merchants ask complex questions requiring multiple data sources, Sidekick coordinates the analysis behind the scenes in real-time. Shopify leverages Claude's advanced reasoning capabilities and Google Cloud's infrastructure to support Sidekick because merchants need around-the-clock access to expert-level guidance.

Results and impact:

- **Millions of merchants** receive expert guidance 24/7
- Entrepreneurs reach **first sale in days rather than weeks**
- Non-developers build internal tools in **minutes rather than days**

Lovable enables users to ship code 20x faster than manual development

Lovable aims to democratize software development by enabling anyone to code. But Lovable faced a technical challenge that had long stumped the industry: how do you help non-coders generate functional, well-architected software? How do you lower the barrier to creating software?

The challenge: Software development timelines blocking idea validation

Traditional software development remains inaccessible to most people, limiting who can ship digital products quickly. Lovable needed an AI solution that could transform conversational interactions into functional, well-architected software that non-technical users could build and technical users could accelerate.

The solution: Conversational interface delivers production-ready code 20x faster

After evaluating commercial and open-source AI models through rigorous quantitative evaluation, Lovable chose Claude for its superior code generation capabilities. By combining Claude's capabilities with their innovations in prompt engineering and system design, Lovable created a platform that delivers functional, production-ready code.

Results and impact:

- **20x faster development** than writing code manually
- **\$40 million ARR** within six months of launch
- **1 million+ active users** monthly building software products

Startups

The barrier to building software has collapsed. Startups are being built in hours. Non-technical founders can prototype working products through conversations with AI agents. Teams of one can ship what once required significant engineering resources.

Replit enables deployment in minutes from any device

Replit is democratizing coding by making it faster and easier to build software. Their mission: empower the next billion software creators—not just developers, but anyone armed with an inspiring product idea. Despite recent AI advances, their mission to democratize access to the digital economy required an easier way to deploy applications for those lacking programming skills.

The challenge: Removing the barrier of learning to code

The time and resources it takes to learn programming remains a significant barrier, even with recent AI advancements. Non-developers needed a solution that could automate complex technical tasks while maintaining professional output.

The solution: Build and deploy applications in minutes without writing code

After evaluating multiple AI models, Replit selected Claude on Google Cloud's Vertex AI for its superior code generation capabilities and seamless cloud integration. Claude demonstrated unique strengths in creating interconnected files that form working applications and editing multiple files simultaneously while maintaining context across the entire codebase.

Running Claude on Vertex AI provided the enterprise-grade security and scalability Replit needed for their global user base. The platform maintains enterprise-grade performance while supporting development from any device, allowing users to go from idea to application in minutes without prior coding experience.

Results and impact:

- Users build and **deploy apps in minutes**
- **Tens of thousands of applications** running on Google Cloud Run
- Non-technical teams across sales, operations, and marketing now building apps

Parcha reduces customer due diligence from 3 months to 5 minutes

Parcha helps the world's leading fintechs and banks scale their customer due diligence processes with AI. The company serves clients including Airwallex, Pipe, Flutterwave, and Alloy—navigating complex regulatory requirements while maintaining rapid growth..

Every financial institution approaches compliance differently. Each monitors transactions and conducts customer due diligence their own way, creating a paradox for Parcha: how do you build flexible AI tools when every customer has unique requirements?

The challenge: Rigid workflows couldn't adapt across institutions

Parcha spent two years developing their own workflow engine but struggled to make it work flexibly across their customer base. Traditional approaches led the engineering team to build rigid workflows for specific customers that were nearly impossible to adapt for others.

Beyond these enterprise workflow challenges, compliance analysts at every institution spend hours conducting open-source research on businesses and individuals—work that's time-consuming, unauditable, and difficult to standardize.

The solution: Adaptive AI system compresses 3-month workflows into 5 minutes

Parcha's tools undergo some of the most rigorous compliance reviews in financial services. Their bank-grade compliance tools are built and certified using Claude models, passing some of the most challenging model governance evaluations in the industry.

After two years building in-house agents with Claude models, Parcha selected the Claude Agent SDK based on proven results and deep familiarity with its capabilities. The Agent SDK allowed Parcha's existing product to gain new flexibility—rather than building separate workflows for each customer's unique requirements, the Agent SDK enabled their product to adapt dynamically.

With the Agent SDK, Parcha developed a fully agentic open-source intelligence research product in just two weeks.

Results and impact:

- Customer due diligence workflows reduced from **3 months to 5 minutes**
- AI tools pass rigorous **bank compliance certifications**
- **\$1 million in new bookings** from platform licensing model

Financial Services

Few industries face more stringent regulatory requirements than financial services, making it the ultimate test for enterprise AI. When banks and investment firms deploy agents at scale, they're proving AI can meet the highest standards for accuracy, auditability, and compliance.

NBIM manages \$1.7 trillion with 20% time savings

Norges Bank Investment Management (NBIM) manages Norway's Government Pension Fund Global, with the mandate to transform Norway's oil revenues into long-term financial wealth for future generations. With \$1.7 trillion in assets, NBIM is one of the world's largest sovereign funds.

Managing \$1.7 trillion across global stock markets, bonds, real estate, and renewable energy infrastructure requires processing massive amounts of information daily. NBIM's teams analyze research reports, market data, regulatory filings, and multilingual news to make investment decisions with enormous consequences.

The challenge: Finding an AI that the entire team could use

The organization struggled to find an AI-powered solution sophisticated enough for institutional investment analysis yet simple enough that everyone—from highly technical portfolio managers to non-technical compliance specialists—could use it.

Complex ESG reporting expectations compounded the challenge, requiring nuanced analysis and clear reporting for each of the 9,000 companies they invest in. They needed AI that was advanced enough for serious investment research and ESG analysis, straightforward enough for their whole team to use, and able to adhere to strict data protection and compliance regulations.

The solution: Human-supervised AI saves analysts 20% of their time weekly

NBIM evaluated multiple LLMs and built human-in-the-loop evaluations to test model capabilities in finance-specific domains. Claude consistently performed best on analysis and reasoning tasks, and in maintaining context over long analytical sessions involving multiple documents. For analysts managing large portfolios, these capabilities translate to 20% time savings every week.

Anthropic's focus on responsible AI development matched NBIM's values and transparency expectations as a public entity. The partnership approach was equally important—the two companies regularly consult on financial services capabilities, test new features, and exchange feedback on safety and enterprise controls.

Results and impact:

- **Weekly 20% time savings** across all departments
- **600+ active users** within two months
- **300 daily Claude Code users**, now most-used AI tool on engineering team

N26 achieves 70% automation across targeted processes in one year

N26 is building the bank the world loves to use, serving diverse, digital-native customers across 24 European markets. As a fully licensed bank, N26 provides a 100% digital banking experience powered by advanced digital, data, and AI technologies.

As N26 rapidly scaled its customer base across multiple European markets, the company faced a critical challenge: ensuring service quality could keep pace with growth. Many essential processes remained manual, document-heavy, and time-consuming.

The challenge: Manual processes couldn't keep pace with rapid customer growth

The rising volume of customer interactions created bottlenecks that directly affected both customer experience and internal efficiency. Teams spent increasing time on document classification, information extraction, translation, and drafting materials based on complex multilingual sources.

The solution: 15+ integrated AI applications achieve 70% automation in one year

N26 selected Claude for its advanced reasoning and multimodal capabilities, which proved essential for handling the sophisticated decision-making required in financial processes. Claude's availability through AWS Bedrock in Europe provided the regulatory compliance, security, and scalability that N26 required as a fully licensed bank.

Since beginning work with Claude in 2024, N26 has integrated the AI assistant into more than 15 internal applications. The deployments span the full customer service lifecycle: a customer-facing virtual assistant providing instant support in five languages around the clock, chargeback request processing that translates documentation and analyzes complex claims, and financial crime analysis where Claude assists analysts by synthesizing customer data and auto-drafting investigation reports.

Results and impact:

- **70% automation** across targeted processes within one year
- **15+ AI applications** serving customer service teams and fraud analysts
- **1-2 weeks** from implementation to testing while meeting regulatory standards

Cybersecurity

Cybersecurity teams face an asymmetric problem: attackers only need to be right once, defenders need to be right every time. AI agents are enabling security teams to analyze threats and respond to incidents at machine speed while maintaining the strategic judgment that only humans can provide.

eSentire compresses threat analysis from 5 hours to 7 minutes

eSentire protects critical infrastructure organizations in 80+ countries as the authority in managed detection and response (MDR). Looking to expand into new markets and protect more customers, the company's security analysts were spending 5 hours on each threat investigation. Their expertise couldn't scale at the pace they needed.

The challenge: 5-hour expert investigations needed to happen at scale

With their Atlas Platform successfully delivering complete threat resolution, the company set their sights on expanding to new markets while deepening customer engagement through enhanced security operations.

eSentire needed to deliver expert-level investigation precision at scale while enhancing transparency of threat resolution outcomes. The goal was enabling existing security experts to amplify their capabilities, protecting more customers while deepening threat analysis in every investigation. What previously required 5 hours of expert analysis needed to happen much faster without sacrificing quality or thoroughness.

The solution: Autonomous AI threat analysis achieves 95% expert alignment in minutes

eSentire evaluated multiple AI models across real-world security scenarios. Claude provided the highest performance for complex security reasoning, with continuous improvements across model versions. Claude's agentic capabilities excelled at orchestrating multi-tool workflows while maintaining investigative coherence—essential for their MDR approach, while its threat analysis aligned

with their most senior security experts 95% of the time while delivering answers in 7 minutes instead of 5 hours.

Amazon Bedrock provided the enterprise-grade security and infrastructure eSentire needed for sensitive intelligence. Claude's intelligent tool selection mirrored expert analyst approaches to complex threat analysis, synthesizing evidence from multiple sources, correlating disparate security events, and incorporating findings into comprehensive conclusions.

eSentire conducted rigorous validation using 1,000 real-world investigations, comparing Claude's decisions against their most senior SOC (Security Operations Center) experts.

Results and impact:

- Expert security analysis compressed from **5 hours to 7 minutes** with **95% alignment**
- **99.3%** threat suppression across critical infrastructure deployments
- **\$1 million+** in new bookings from platform licensing

Palo Alto Networks accelerates junior developer integration by 70%

Palo Alto Networks, the world's largest cybersecurity company, protects organizations globally with comprehensive security solutions. As threats evolve faster, the engineering team saw an opportunity to use AI to help their products stay ahead of competitors and bad actors.

The challenge: Accelerating development without compromising security

The engineering team mapped their software delivery process to identify where developers spent time and where errors mostly occurred. Their analysis revealed that while junior developers spent 30-35% of their time in initial development, this phase was also where the most critical issues emerged. And rushing developers through onboarding to ship features faster than competitors could open them up to unintentionally shipping vulnerabilities, too.

They needed to compress the months-long path it often took to get their engineers up-to-speed on their codebase and processes without creating the security debt that comes from moving too fast. Palo Alto Networks saw a powerful opportunity: using generative AI to stay ahead of both competitors and bad actors in the rapidly evolving cybersecurity landscape.

The solution: AI-powered development cuts junior developer onboarding time by 70%

After evaluating multiple AI solutions, the engineering team chose Claude on Google Cloud's Vertex AI for several key reasons. Claude consistently demonstrated superior performance in coding tasks while maintaining high accuracy and security standards.

Anthropic prioritized safety and security a lot more than other LLM providers, discussing security and safety implications in every meeting, a critical consideration for a cybersecurity leader.

Using Google Cloud's Vertex AI provided significant advantages, including granular usage-based pricing with flexible commitment periods and seamless SDK integration for provisioned throughput.

Results and impact:

- Junior developers complete complex integrations **70% faster**
- New developers contribute meaningfully in **weeks instead of months**
- **20-30% increase** in feature development velocity

Legal

Legal and professional services built their business models on expert knowledge and billable hours. AI agents are transforming this equation, making decades of expertise instantly accessible while freeing professionals to focus on client relationships rather than research and document review.

Thomson Reuters delivers comprehensive legal analysis in minutes

Thomson Reuters is a global content and technology company serving legal, tax, accounting, and compliance professionals. With 3,000 domain experts and over 150 years of authoritative content, the company provides trusted intelligence to professionals who need accurate information to advise their clients.

As a result of their size and breadth of expertise, they faced a knowledge access crisis that directly impacted their users' ability to serve their clients. A lawyer researching a case precedent, for example, would have to manually search through thousands of documents, billing hours while hunting for relevant materials.

The challenge: 150 years of legal expertise trapped in formats lawyers couldn't access quickly

Thomson Reuters needed an AI solution to leverage their vast professional knowledge base while meeting strict accuracy and reliability requirements. The

risk of providing wrong advice is substantial, requiring extraordinarily high quality thresholds.

They needed technology to maintain their high standards while making expert knowledge more accessible. In professional services, response speed can directly determine who wins clients. But accuracy can't be sacrificed, as wrong advice can expose clients to regulatory liability and penalties.

The solution: AI platform delivers comprehensive legal analysis in minutes instead of hours

Thomson Reuters uses Claude in Amazon Bedrock as part of its strategy to power their AI platform, CoCounsel, helping legal and tax professionals synthesize expert knowledge and deliver comprehensive advice to clients.

After extensive evaluation against automated and human expert benchmarks, Thomson Reuters selected Claude. Their choice to deploy Claude in Amazon Bedrock emerged from three key factors: Anthropic's focus on safety aligning with their core values, Claude consistently meeting rigorous quality standards, and Amazon Bedrock providing significant enterprise advantages for rapid testing and deployment while maintaining strict security standards.

Results and impact:

- **Minutes to search** 3,000+ domain experts and 150 years of case law
- Thomson Reuters tests AI models **the day they're released** via Amazon Bedrock

A | A | A | A | A - A - A - A | A | A |
A | A | A - A - A | A | A | A | A | A | A |
A | A - A - A | A | A | A | A - A | A | A |
A | A - A | A | A | A | A | A | A - A | A | A |
A | A | A | A | A | A | A | A - A | A | A |
A | A | A | A | A | A | A - A | A | A | A |
A | A | A | A | A - A - A | A | A | A | A | A |
A | A | A | A | A - A - A | A | A | A | A | A |
A | A | A | A | A - A - A | A | A | A | A | A |

Outlooks and perspectives

Outlooks and perspectives

Leading services providers see 2026 to mark AI agents' transition from pilots to production systems. [**Anthropic's 2025 Economic Index**](#), which analyzed over 3.5 million Claude conversations, shows that shift is already underway—and reveals where agentic adoption succeeds and stalls.



Accenture: Conversational interfaces will define enterprise AI

In 2026, AI agents will become essential enterprise tools, handling complex workflows and team collaboration at scale. The key enabler: conversational interfaces that let users interact with AI through natural dialogue rather than rigid commands.

Standalone language models struggle with complex tasks because they lack external context and can't take action independently. Pairing them with conversational interfaces unlocks capabilities like coding through dialogue.

The technical challenge is resilience. Traditional stateless systems are optimized for scalability but lose context when disrupted. AI interactions need to tolerate interruptions—network glitches, user breaks, system restarts—while maintaining continuity, similar to how human conversations naturally resume after disruption.

Accenture's Center for Advanced AI is building infrastructure based on Model Context Protocol (MCP) to enable stateful, resilient AI communication with features like resumability, redelivery, and direct LLM sampling. The goal: agents that work seamlessly across platforms while preserving progress and context.

"2026 will separate enterprises that deployed AI agents from those that transformed around them. The ROI ceiling isn't set by the technology—it's set by the willingness to redistribute authority, redesign workflows, and trust intelligent systems with consequential decisions. Companies treating this as a technology implementation challenge will see incremental gains. Those recognizing it as a reinvention imperative will create compounding advantages that become impossible to replicate."

Alex Holt, Vice Chair and Global Strategy Leader, Tech, Media & Comms, Accenture



Boston Consulting Group (BCG): 2026 will be the year agents drive real business impact

With low-value experimentation behind us, 2026 will mark the shift to agents driving measurable impact on business results beyond coding. Early successes have come from constrained agents—systems designed for specific, bounded tasks. Autonomous agents promise more scalability and flexibility, with the capacity to change how companies operate if built well.

This requires multiple changes happening simultaneously. Workers at all levels will need to reskill to collaborate with agents. Quality controls and managerial systems need to evolve to handle non-human work product. More agent-ready infrastructure will emerge—like remote MCP servers—making broader ecosystems accessible to agents. The first reliable agent-to-agent workflows will appear in consumer contexts before enterprise deployment. This increased connectivity will require companies to streamline and integrate their data, design their own agent ecosystems, and implement change programs spanning people, technology, and AI.

"If 2025 was the 'year of agents', 2026 will be the year we start to put them to work. The underlying AI models are racing ahead, but in the enterprise—legacy technology and processes lag behind. We find clients succeed and realize P&L impact faster when they focus on transforming their systems end-to-end with agents at the center, rather than as a tack-on to legacy processes. AI transformation must serve a business strategy, with the business challenges to solve as a North Star."

Tom Martin, Director, AI Platforms, Boston Consulting Group



Deloitte: Workforce transformation will determine who wins

Organizations will pivot from exploring agentic AI to scaling it in 2026 by redesigning workflows and entire business units. The hybrid human-machine workforce is coming into focus, but technical deployment is only half the challenge. The deeper risk is failing to sustain behavior change after workflows are transformed.

For enterprise AI programs to succeed, employees need to adopt new processes without regressing to old behaviors. This requires clear communication about how working with agents improves outcomes and enhances their workplace experience. It goes beyond training people to use new tools—organizations need to upskill their workforce and incentivize participation in transformed ways of working. Agent deployment and human behavior change need to happen in tandem.

"Even as organizations begin to unify their IT estate to capitalize on the capabilities agentic AI enables, they will also need to unify their workforce behind the transformation of work, generating buy-in and enthusiasm rather than skepticism and reluctance."

Jim Rowan, Head of AI, Deloitte US

Perspectives from Anthropic's 2025 Economic Index

Anthropic's [2025 Economic Index](#) analyzed over 3.5 million anonymized Claude conversations to understand how AI is being used across industries. The findings reveal a clear pattern: enterprises are moving beyond experimentation toward systematic deployment, with usage concentrated in areas where AI capabilities are strongest and organizational barriers are lowest.

Enterprises are delegating, not collaborating

77% of business API usage shows automation patterns, meaning companies are handing off complete tasks to AI rather than using it as a collaborative assistant. This is significantly higher than consumer usage, which hovers around 50%. Enterprises are embedding AI into workflows as a workhorse, not necessarily as a thought partner. This aligns with our survey findings: 97% of respondents expect increased efficiency gains from their agentic deployments over the next 12 months.

Capability matters more than cost

The most expensive tasks have the highest usage rates. Businesses are deploying where model capabilities are strong and where automation creates real economic value. For technical decision makers, this suggests the ROI calculation should focus on business outcomes, not token costs. Complex code generation, multi-step research synthesis, and detailed document analysis all require more compute but deliver outsized returns when done well. Our survey with Material found similar sentiment: 96% of respondents felt optimistic about the business impact of AI agents at their company.

Context is the real bottleneck

Complex tasks require disproportionately more context to execute well. There's a stable relationship across tasks: every 1% increase in input context length is associated with a 0.38% increase in output quality and length. For some organizations, costly data modernization and investments to surface contextual information may be the primary bottleneck for AI adoption. Companies with fragmented or siloed data will struggle to unlock sophisticated AI use cases.

Usage is concentrated in predictable places

Nearly half of Claude's API traffic—44%—maps to computer and mathematical tasks. AI deployment succeeds where model capabilities are strong, deployment barriers are low, and employee adoption is quick. Tasks requiring regulatory approval, dispersed tacit knowledge, or capabilities AI can't handle see minimal adoption. Our survey reflects this pattern: nearly all organizations are adopting coding agents, followed by similarly quantitative tasks including supply chain optimization and financial planning.

The shift from augmentation to automation is accelerating

Over eight months, directive conversations—where users delegate complete tasks—jumped from 27% to 39%. This marks the first time automation usage exceeds augmentation. Whether this shift is driven by improving model capabilities or users learning to trust delegation, the direction is clear: enterprises are moving toward full task handoff. Our survey found the same: more than half of businesses already use AI agents to handle complex, multi-step workflows like customer service resolution and employee onboarding. These workflows often involve sensitive customer data, proprietary business logic, and access to internal systems—making trusted, secure AI technology essential for production deployment.

Geographic maturity predicts usage diversity

High-adoption regions show diverse AI applications across education, science, and business operations. Lower-adoption countries concentrate over 50% of usage on coding alone. As organizations mature in AI adoption, they tend to diversify beyond their initial use cases—something worth planning for.

The strategic takeaway

The biggest barriers to enterprise AI value aren't model capabilities or costs—they're organizational readiness, especially around data accessibility and context aggregation. Companies that invest in making their internal knowledge accessible to AI systems will be better positioned to capture value from increasingly capable models.

A I A I A I A I P I P I P I P I A I A I
A I A I A I P I P I P I P I A I A I A I A I
A I A I P I P I P I P I A I A I A I A I A I A I
A I P I P I P I P I A I A I A I A I A I A I
P I P I P I P I A I A I A I A I A I A I
P I P I P I A I A I A I A I A I A I A I A I
P I A I A I A I A I A I A I A I A I A I A I
A I A I A I A I A I A I A I A I A I A I A I
A I A I A I A I A I A I A I A I A I A I A I
A I A I A I A I A I A I A I A I A I A I A I

Additional
resources

Additional resources

Setting your AI strategy for 2026? Check out other resources from our team:

Guides

- **The Enterprise AI Transformation Guide:** Learn how to accelerate the deployment and adoption of AI at your organization.
- **Scaling Agentic Coding Across Your Organization:** Learn best practices and guidance from our Applied AI and engineering teams about how to drive agentic coding adoption at your company.
- **Building Effective AI Agents:** Architecture Patterns and Implementation Frameworks: Learn how technical teams at Anthropic and other organizations are building agentic workflows, single-agent systems, and multi-agent solutions to power enterprise use cases. Includes real-world agentic architectures you can deploy today.

Blog articles

- **Building AI agents for startups:** Learn how leading startups like eSentire and Gradient Labs build and deploy AI to accelerate the software development lifecycle, scale operations with lean teams, and even prototype new products.
- **Building AI agents in healthcare and life sciences:** Learn how leading healthcare and life sciences organizations like Novo Nordisk and Pfizer build AI agents that expedite drug discovery and clinical trials, streamline operations, and ensure regulatory compliance at scale.
- **Building AI agents for financial services:** Learn how leading financial services institutions like NBIM and Intuit Turbo Tax build AI agents that combat fraud, improve customer experiences, and generate real-time market analysis.

Start building with Claude

[Contact our Sales team](#) to build agents that handle complex workflows on the Claude Developer Platform or integrate Claude directly into your development environment with Claude Code.

