




AI agent trends 2026

Five shifts that will
redefine roles, workflows,
and business value in 2026.



About this report

This report provides key insights for business leaders to shape their AI agent strategy for 2026 and beyond. Within each trend, you will find real-life examples, technical resources, and customer stories to share with your teams for deeper learning.

These trends were identified using a blend of qualitative and quantitative data, including internal Google Cloud and Google DeepMind interviews with AI leaders, customer case studies, and insights from [The ROI of AI 2025](#) report (based on a global survey of 3,466 enterprise decision makers), with analysis using NotebookLM and Google AI Studio.

AI agents are redefining business value in 2026

The ceiling for human achievement has been lifted.

While headlines focus on the distant future of Artificial General Intelligence (AGI), the decisive shift for business is happening right now.

That shift is **agentic AI**. This is AI that moves beyond answering questions to understanding a goal, making a plan, and taking actions across applications to achieve it with extensive human guidance and oversight.

Due to this capability, AI is one of the first technologies that applies to every single person, in both our personal and professional lives. Its power lies in its ability to augment human capacity with better recall, faster data processing, and enhanced reasoning across the back office, front office, and corner office. Achieving this, however, depends on ensuring widespread access, necessary skills development, and broad participation for everyone.


In this report, we explore five key AI agent trends shaping business in 2026. Unlocking the value of these trends requires more than simply adopting new tools. It also demands that leaders question old assumptions and drive the cultural change necessary to thrive in this new, agentic AI era.

What are AI agents?

Agents are systems that combine the intelligence of advanced AI models with access to tools so they can take actions on your behalf, under your control.

| I/O, May 2025

| Keynote by Sundar Pichai, CEO, Google



“ AI agents are the leap from being an ‘add-on’ approach to being an ‘AI-first’ process. It’s a fundamental change in workflow, a new way to work that will require a profound shift in mindset and corporate culture.”



Oliver Parker

Vice President, Global GTM for Generative AI,
Google Cloud

5 AI trends shaping 2026


1 Agents for every employee

Empowering individuals to achieve peak productivity




2 Agents for every workflow

Running your business with grounded agentic systems




3 Agents for your customers

Delighting customers with concierge-like experiences



4 Agents for security

Advancing security from alerts to action



5 Agents for scale

Upskilling talent will be the ultimate driver of business value





Agents for
every employee

Empowering individuals to achieve peak productivity


A large, abstract, three-dimensional mesh structure in shades of red and purple, resembling a stylized human figure or a complex network, positioned on the right side of the page.



The most significant business shift of 2026 isn't just about efficiency; it's a fundamental, employee-centric transformation.

This new agentic model is designed to expand the potential of every individual, turning them into the primary engine for innovation and growth.

This change stems from a behavioral shift in the human-computer interface, moving from instruction-based computing (e.g., analyzing a spreadsheet, developing code) to intent-based computing. In 2026, employees will be increasingly able to state a desired outcome, and the computer—using LLMs and agents—determines how to deliver it.



52%

of executives in gen AI-using organizations have AI agents in production,¹ deploying them across a wide range of use cases:²

Of those:

- 49% use agents for customer service
- 46% use agents for marketing or security operations
- 45% use agents for tech support
- 43% use agents for product innovation or productivity and research

¹ Google Cloud, [The ROI of AI](#), 2025 (Total global: n=3466; Questions: (1) How are AI agents leveraged across the enterprise? (2) How many AI agents does your company currently have deployed in production across your organization?)
² Google Cloud, [The ROI of AI](#), 2025 (Executives whose organization is leveraging agentic AI: n=1814; Question: What use cases has your company deployed AI agents for?)



“By 2026, agents will manage complex, multi-step workflows across systems. A key responsibility of employees will be to set the strategy and oversee the system of agents responsible for tasks, such as invoicing and contracting.”



**Saurabh
Tiwary**

VP, General Manager,
Cloud AI, Google Cloud



“

This surge in AI [agent] adoption reflects a fundamental mindset shift to recognizing it as a productivity instrument available 24/7. At TELUS, we’ve seen this firsthand—over 57,000 team members regularly use AI and save 40 minutes per AI interaction.”

Jaime Tatis

Chief AI Officer, TELUS



A new integrated working model

In this new model, every employee—from an entry-level analyst to a senior vice president—becomes a human supervisor of agents.

Their primary job is no longer to perform every mundane task personally, but rather to orchestrate a team of specialized AI agents to achieve a goal. This model is about more than just delegation; it's about augmentation. The real power comes from giving every employee agents grounded in the company's own enterprise context—its internal systems, knowledge bases, customer data, and past work—to elevate the impact of their efforts.

The employee's core function becomes providing strategic direction. Their new responsibilities are to:

Delegate mundane or repetitive tasks

Identify which tasks are best suited for an agent and assign them.

Set goals

Clearly define the desired outcome for the agent.

Outline strategy

Use their human judgment to guide the agents and make the final, nuanced decisions that AI can't.

Verify quality

Act as the final checkpoint for quality, accuracy, and tone.

What is grounding in AI?

It is the process of anchoring an AI model's responses to a specific, verifiable set of facts—its “ground truth.” For an enterprise, this ground truth is its own internal data.

[Learn more](#)



As roles shift to agent management, enabling employees is vital, especially in regions like Japan that rely on system integrators (SIs). Democratization via tools like Gemini Enterprise allows knowledge workers to build agents, improving productivity and elevating the SI partnerships to focus on complex, long-term initiatives.”



**Hiroyuki
Koike**

Managing Director,
Customer Engineering,
Japan, Google Cloud

How it works

The 10x marketing manager

A marketing manager’s job used to be a constant scramble of drafting posts, pulling data, and watching competitors. In 2026, they can orchestrate a system of specialized AI agents to achieve their goals, rather than performing every task personally.

With agents focusing on specific tasks, the marketing manager can multiply their output by focusing on high-impact brand storytelling and strategic campaign development.

 Ask anything, search your data, @mention or /tools







Enable all connectors

 Google Search

 Calendar

 Drive

 Gmail

 Manage connectors

An example of a Gemini Enterprise home page

Google Cloud

11

The marketing manager's specialized agent system

Their new role involves orchestrating five specialized agents:

Data agent

This agent can sift through millions of structured and unstructured data points to find actionable patterns in market trends.

Analyst agent

Its job is to monitor market trends, competitor announcements, and social media sentiment 24/7. It delivers a one-page report of key insights to the marketing manager's inbox every morning.

Content agent

Given a strategic theme for the week, the agent drafts copy for several social media posts and a blog article in the company's brand voice, then shares it with the marketing manager for review.

Creative agent

The marketing manager supplies the agent with a marketing strategy, creative guidelines, and social copy. The agent then generates images and videos to accompany those social posts and shares them with the marketing manager.

Reporting agent


This agent connects to the company's analytics platforms. With the marketing manager's guidance and oversight, this agent pulls and analyzes weekly campaign data and delivers a one-page summary of key insights every Friday.



Data tip

Teams can use data agents as an intelligence multiplier.

[Learn more](#)






Suzano, the world’s largest pulp manufacturer and a leader in sustainable bioeconomics, worked with Google Cloud and Sauter to develop an AI agent built with Gemini Pro to translate natural language questions into SQL code to query SAP Materials data on BigQuery. This has resulted in a **95% reduction in the time** required for queries among the 50,000 employees using the data.



Ready to 10x your team?

Gemini Enterprise enables employees to build and manage their own specialized AI agents.

Try now



“ There is a common misconception that agents act without control. Humans will remain the orchestrators and final decision-makers. Agents will function as powerful assistants to augment human-centric workflows.

Take the media and entertainment industry, where agents can help in the understanding of vast amounts of complex content and data—and human creative and strategic thinking will remain critical to deciding which stories to tell, how to tell the stories, how to fund and distribute and monetize and market these stories, compliance, rights management, and of course, where not to invest time and effort.”



**Albert
Lai**

Director, Telco, Media & Technology,
Global Strategic Industries,
Google Cloud



Agents for
every workflow

Running your business with grounded agentic systems





An agentic system is a digital assembly line—a human-guided, multi-step workflow that orchestrates multiple agents to run a business process end to end.

The true value in 2026 begins with augmenting the capabilities of individuals and teams, and continues to grow by making the entire business run more intelligently and efficiently, 24/7, at scale.



88%

of agentic AI early adopters are now seeing a positive ROI on at least one gen AI use case³

“ AI is driving a generational refactoring of the enterprise—the core workflows and the entire technology stack. AI agents will transform complex, multi-step processes like procurement, security operations and customer support—shifting the human roles to focus on high-value, strategic orchestration across the business.”



Francis
deSouza

COO and President,
Security Products,
Google Cloud

³ Google Cloud, [The ROI of AI](#), 2025 (Agentic AI early adopters: n=460; Question text: In what timeframe do you expect gen AI to deliver return on investment (ROI) to the following areas of your business?)



“

The real value will lie in using agentic workflows to integrate siloed functions, such as network operations, field services, and customer contact centers in telecommunications. Agents could autonomously remediate network anomalies, proactively open a ticket with the field service systems, and alert contact centers to inform customers of a technician dispatch, all in one integrated sequence.”



Angelo Libertucci

Director, Telecom,
Global Strategic Industries, Google Cloud


Digital assembly line: Orchestrating agentic systems

A digital assembly line is made possible by the Agent2Agent (A2A) protocol.

This open standard enables seamless integration and orchestration between AI agents, allowing them to work together even if they are from different developers, built on different frameworks, or owned by different organizations.

While LLMs are the “brains” of these agents, they have two major limitations: their knowledge is frozen at the time of their training, and they can’t interact with the outside world to access real-time data or perform actions.

The Model Context Protocol (MCP) solves this. It creates a standardized, two-way connection for AI applications, allowing LLMs to easily connect with various data sources and tools, such as managed databases (e.g., [Cloud SQL](#), [Spanner](#)) and data platforms (e.g., [BigQuery](#)).





Salesforce is working with Google Cloud to create AI agents that work across both platforms using the newly launched Agent2Agent (A2A) open protocol, a leap forward in building an open, interoperable foundation for agentic enterprises.



Elanco, a global leader in animal health, uses Gemini models within its Elanco.ai platform to automatically sort, extract key insights from, compare, and restructure information from over 2,500 unstructured policy and procedure documents per manufacturing site. The AI agent improves accuracy and consistency, reducing the risk of outdated or conflicting information that **could cost up to \$1.3 million** in productivity impact at large sites.



Data tip

Discover how your data teams can ground AI in your business data.

[Get started](#)



Over the next few years, we will see significant growth in Agent2Agent operations implemented across industries. For example, an AI agent from a media company could connect to a retailer's agent to showcase the details and pricing on a specific product shown in streamed or broadcast content. Similarly, AI agents at hospitals could work directly with labs or insurance agents, provided the patient grants permission to handle their sensitive health data.”



**Marcel
Silva**

AI GTM & Sales LATAM,
Google Cloud


How it works

Agentic ecommerce in action

Today's payment systems assume a human is directly initiating the purchase.

This poses a fundamental challenge for security: What happens when a non-human entity (the agent) is making the final transaction decision, with pre-approval from a human?

An agent initiating a payment with a human's oversight and guidance under a new framework like Google Agent Payments Protocol (AP2) breaks this assumption. This raises critical questions about how to prove user-given authority for a purchase, how a merchant can be sure an agent's request is accurate and not a hallucination, and who is ultimately accountable in case of fraud.



For example, a customer discovers a winter jacket they want is unavailable in a specific color. They can then tell their agent:

“Purchase this jacket when it becomes available in black. Don’t purchase it if the price is more than \$100.”

The agent then monitors prices and availability and, with human pre-approval, executes a secure purchase the moment that specific variant is found, capturing a high-intent sale that would have otherwise been lost.



Try Vertex AI

Get started



PayPal is creating agentic shopping and commerce experiences through adoption of industry leading protocols such as Google’s Agent Payments Protocol (AP2), a secure, open, scalable solution paving the way for the future of agentic commerce.





“

In 2026, agentic ecosystems will scale from pilots to production in financial services. For example, we'll see multi-step agentic compliance systems that can monitor regulatory changes, identify impacted policies, update internal workflows, and create a complete audit chain, making the process more efficient.”




Toby Brown

Managing Director, Financial Services,
Global Strategic Industries, Google Cloud



Agents for
your customers

Delighting customers with concierge-like experiences

A large, abstract, three-dimensional shape composed of many small, glowing green and blue dots, resembling a digital or data-based structure. It is positioned in the lower right quadrant of the page, partially overlapping the main text.



For the last decade, customer service automation meant pre-programmed chatbots answering simple questions and deflecting support tickets.

They were efficient, but they lacked the ability to understand more nuanced and complex questions.

With advances in LLMs and A2A, 2026 will deliver more helpful concierge-style agents. These AI agents will connect enterprises and customers by remembering preferences and past conversations to offer truly one-to-one experiences.



49%

of executives at organizations with AI agents in production report adopting agents for customer service and experience⁴

“Current call center automation systems require callers to go through scripted options or pre-programmed chats, often requiring them to repeat ‘operator!’ to reach a human. Agents allow for quicker, more natural interaction by letting customers speak and provide context. This return to verbal communication will be a reality in the next 1–3 years.”



Paul
Tepfenhart

Director, Retail & Consumer,
Global Strategic Industries, Google Cloud

⁴ Google Cloud, [The ROI of AI](#), 2025 (Executives whose organization is leveraging agentic AI: n=1814; Question text: What use cases has your company deployed AI agents for?)



Personalization at scale

Your customer no longer has to start every conversation by proving who they are or re-explaining their problem.

Chatbot

“Please enter your 12-digit order number.”


The difference isn't just the AI, it's the data. The agentic concierge can succeed because it is grounded in an enterprise context you choose to share for a particular customer—from purchase history in the CRM to package tracking in the logistics database.

Agentic concierge

“Hi, Elizaveta. I see you're calling about the blue sweater you bought last week. Our system shows it was just delivered. Are you calling to start a return or an exchange?”



Home Depot built Magic Apron, an AI agent that offers expert guidance 24/7, providing detailed how-to instructions, product recommendations, and review summaries to make home improvement easier.



“ Personalization extends beyond consumer experiences. On the manufacturing floor, for example, agentic systems could offer personalized advice to managers. If the second shift underperforms the first, the system could inspect multiple machine criteria and suggest solutions like offering more training or recommending optimal machine set points.”



**Praveen
Rao**

Director, Manufacturing,
Global Strategic Industries,
Google Cloud



How it works

Insightful and always-on help

Customer service is defined by its helpfulness.

An agentic concierge doesn't wait for a complaint. It monitors systems for triggers and resolves problems using real-time data to provide insights and take actions with human guidance and oversight.



Try Gemini Enterprise for Customer Experience

Get started

Vertex AI Search

Agentic search with accurate, instantaneous answers

Shopping agent

Digital concierge with complex reasoning to execute consented actions

Food Ordering agent

Interactive multilingual ordering across channels

CX Agent Studio

Personalized and always-on multimodal customer engagement agents

Agent Assist

In-the-moment coaching and assistance for faster issue resolution

CX Insights

Insights to improve customer service operations performance and efficiency



Building trust through proactive customer service.


Imagine this scenario: a logistics agent flags a delivery as “failed” at 3 PM. Instead of waiting for an angry customer to call, the concierge agent could take the following steps:

- 1 Checks the backend and confirms the delivery van broke down.
- 2 Accesses the logistics system to reschedule the delivery for the first-available slot tomorrow morning.
- 3 Logs in to the billing system and applies a \$10 service credit for the inconvenience.
- 4 Notifies the customer via text:
“Hi, we’re so sorry. Your package was delayed due to a vehicle issue. We’ve added a \$10 credit to your account for the trouble. I’m happy to reschedule it for tomorrow between 9–11 AM. Please reply ‘Yes’ to confirm the new time or respond with your preferred time.”

The problem is resolved quickly. The agent is designed to execute a “smart handoff” with a full summary for complex or emotionally charged issues; human staff can step in if the case becomes much more complex.



Danfoss, a global manufacturer operating in over 100 countries, uses AI agents from Go Autonomous on Google Cloud to automate email-based order processing. The solution automated 80% of transactional decisions, **reduced average customer response time from 42 hours to near real-time**, and consolidated five systems into a single interface.



“ The greatest promise of agentic AI in 2026 is moving healthcare from a reactionary system into a predictive learning health system. For providers, this means agents integrating data—from imaging and EHRs to claims—to delivering proactive insights directly into the clinician’s workflow. This would allow preemptive risk management across patient populations and the democratization of high-quality healthcare.”



Aashima Gupta

Director, Healthcare,
Global Strategic Industries, Google Cloud

1

2

3

Trend 4

5



Agents
for security

Advancing
security
from alerts
to action





In a modern security operations center (SOC), human analysts face a constant stream of data and alerts, with 82% concerned or very concerned that they may be missing real threats or incidents due to the amount of alerts and data they are faced with.⁵

This “alert fatigue” is the attacker’s greatest advantage; they only need to be right once, while the defender has to be right every single time.

While security orchestration automation and remediation (SOAR) solutions deliver some automation, they may offer only incremental benefits. But with their ability to reason, act, observe, and adjust actions based on new information, AI agents have the potential to help security teams identify and respond to threats more effectively.

Expanded [Secure AI Framework 2.0](#) can help address the rapidly emerging risks posed by autonomous AI agents. Additionally, early results from DeepMind’s research on [CodeMender](#), a new AI-powered agent that improves code security automatically, already demonstrated its ability to find new zero-day vulnerabilities in well-tested software.

In 2026, AI agents will increasingly help with tasks like vulnerability discovery as well as alert triage and investigation.



46%

of executives at organizations with AI agents in production report adopting agents for security operations and cybersecurity⁶

⁵ Forrester Consulting on behalf of Google, [Threat Intelligence Benchmark: Stop Reacting: Start Anticipating](#), July 2025

⁶ Google Cloud, [The ROI of AI](#), 2025 (Executives whose organization is leveraging agentic AI: n=1814; Question text: What use cases has your company deployed AI agents for?)



Today's CISO is laser-focused on achieving the greatest decrease of risk per dollar spent. Agents are essential to this, as they detect and respond faster to enterprise risks. More importantly, they elevate our SOC analysts from tactical responders to strategic defenders."



Jon Ramsey

Vice President & General Manager,
Security, Google Cloud




Specular, an offensive cybersecurity platform, builds AI agents using the Gemini 2.5 Pro model to automate attack surface management and penetration testing. Their platform automates traditional workflows to identify, assess, and remediate cybersecurity, helping enterprises quickly prioritize and respond to threats.

The semi-autonomous security operations cycle

An agentic SOC orchestrates a system of task-based AI agents, each with a specific role, to achieve a common security outcome.

After receiving a security alert, the agentic SOC cycles through a process, engaging various agents:



This dynamic process of evaluating, acting, and re-evaluating enables the system to adapt to a changing security environment in real time, while freeing up time for human analysts to focus on higher-value work.

Multiple SOC agents need common enterprise context and can share the same security data sources (e.g., security telemetry data), regularly communicate, and adapt their actions through technologies like A2A and MCP. Agents should also be trained on continuously evolving real-world insights from security experts.




“

AI is already being used to find zero-days, identify malicious code, and uplift the work of defenders. As threat actors incorporate the technology into their operations, it will be our best tool to meet this new challenge.”



Sandra Joyce

Vice President of Threat Intelligence,
Google Cloud





How it works

Elevating the security analyst

With the addition of agentic systems acting as force multipliers, human analyst roles are expected to change for the better. AI agents can take on the draining, reactive work of “alert-watching,” shifting the human analyst to a more strategic level, engaging in activities such as:

Threat hunting

Using their intuition and experience to guide the agents. Example: “Agent, I have a hunch about this server. Hunt for any unusual outbound data transfers.”

Supervising agents

Fine-tuning the agents’ “rules of engagement” and “performance-reviewing” their automated responses.

Defending

Focusing on long-term security posture, architecting better defenses and anticipating the next wave of attacks, rather than just chasing the current ones.



Explore Google
Security Operations

Get started



torq=

Torq uses agentic AI to automate the entire security operations lifecycle through Socrates, an AI SOC analyst that coordinates specialized agents. Running on Google Cloud’s infrastructure, teams achieve 90% automation of tier-1 analyst tasks auto-remediated without human involvement, **95% decrease** in manual tasks and **10x faster** response times.



AI is transforming the security landscape in both offense and defense, and AI infrastructure including models, data, and agents are dramatically expanding an enterprise's attack surface area. Security professionals must be deeply bilingual in both AI and security, to stay ahead of advanced AI threats from bad actors, and defend against them with sophisticated AI tools."



Francis deSouza

COO and President, Security Products,
Google Cloud

1

2

3

4

Trend 5



Agents
for scale

Upskilling
talent will be
the ultimate
driver of
business
value

A large, abstract, three-dimensional cloud-like structure composed of numerous small, glowing blue and purple spheres. It occupies the right side of the image, partially overlapping the text.



It is tempting to focus on the technology—the models, the platforms, and prompts—but this misses the most critical element: the people.

As AI is evolving, the skills gap is widening and it's getting harder for individuals and organizations to keep up. Skills themselves expire faster than ever: The “half-life” of a professional skill is now four years—and in tech, as short as two years.⁷

Both practitioners and decision-makers see the importance of closing this gap. Skills increase the ability to get hired, get promoted, and grow careers—plus, they have a positive impact on productivity, innovation, and revenue.

“ AI offers an unprecedented opportunity for employees to harness the data and context around them. 2026 will be the year when every employee can go from guessing to knowing—but only if their organizations invest in the skills to make it possible.”



Andrew
Milo

Global Director, Customer Training,
Cloud Learning Services,
Google Cloud

⁷ Forbes, [AI Puts The Squeeze On The Shrinking Half-life Of Skills](#), 2024




What executives are saying:⁸



What employees are saying:⁹



⁸ Google/Ipsos, Cloud Learning Services Market Pulse, fielded Sept.-Nov. 2024 (U.S., U.K., FR, DE, IN, BR, MX, JP, AU/N.Z.; Decision-makers (n=902). Findings reflect the opinion of survey respondents only)
⁹ Google, [Beyond AI Optimism: Five ways to move your business from saving time to sparking innovation](#), 2025





In 2026, the shift of employee scope to include agent management and orchestration will create a skills gap. It's because the expertise to be an 'agent orchestrator' or 'Chief of Staff for AI' simply doesn't exist in the market yet."



Shweta Maniar

Director, Life Sciences,
Global Strategic Industries, Google Cloud



How it works

The 5 pillars of AI learning

To thrive, organizations must move beyond simply buying technology and focus on building an AI-ready workforce. This requires a holistic strategy built on five key pillars.

01

Establish goals

Figure out what is most important and what you can actually measure. For example, perhaps your business goal is 100% AI tool adoption within your organization. So, every team member would use a Gemini Enterprise agent at some point in their workflow—assisting them with increased recall, faster processing, or enhanced reasoning—making the flow better than what they could have done on their own. Make sure your specific business goal aligns to the bigger picture of what the organization needs—and that it's measurable.





02

Secure sponsorship

Gathering a team together ensures that there is constant communication and momentum around your AI initiative. A good place to start is with a team of three primary stakeholders: an **executive sponsor**, the **groundswell lead**, and the **AI accelerator**. Let's break these roles down.

The **executive sponsor** provides the necessary funding, high-level backing, and consistent messaging on AI's importance.

The **groundswell lead** acts as the "AI megaphone" by managing grassroots campaigns, generating excitement, and collecting employee ideas.

The **AI accelerator** is the technical expert responsible for transforming those prioritized ideas into functional solutions.

This structure ensures the program has the support to be funded, the momentum to engage employees, and the technical capability to deliver tangible AI tools.

03

Sustain momentum and reward innovation

To sustain engagement, the program should use a layered strategy centered on interactive platforms and consistent communication. Think of the program as a "digital hub" featuring a gamified idea exchange with a leaderboard to continuously collect and reward AI use cases and new ideas. Peer-to-peer knowledge should be highlighted in a way that showcases successful AI applications across various roles.

Maintain momentum through regular, multichannel communication, including weekly executive emails, content for all team meetings, and a quarterly awards program to recognize and reward top innovators.



04

Integrate AI into daily workflows

Keep the AI education going and tap into your team's collective genius by hosting internal hackathons where small teams compete to develop and pitch innovative AI solutions. Winning concepts could even be integrated into the official program.

Consider hosting Field Days, which are structured, challenge-based events that encourage teams to practice using the new custom AI tools and other innovations in a collaborative, practical setting. These can be both in-person and virtual.

05

Prepare for increasing risks with trusted frameworks

As agent-accelerated cyber security risks become ever more sophisticated, security becomes everyone's responsibility. It's critical to ensure that employees understand their role—including training on what data can (and cannot) be used in AI tools, and how to recognize sophisticated threats like social engineering that uses AI.





“

AI adoption should be an ongoing story, instead of just one chapter in a book. We had found that 96% of team members reported that they had increased confidence using AI tools. And 96% committed to applying these tools in their work. Our Google Skills training program has actually doubled in impact from February all the way till September 2025.”

Gina Narvaez

Director of Digital Talent Team, TELUS



Try Google Skills

Find AI training courses and lessons on a wide range of topics.

[Get started](#)



“

When working with AI, executive boards and leaders must prioritize upskilling their teams in critical thinking and applying ethical judgment in decision making. Since these skills are usually built through experience, a bank, for example, could first roll out its customer support chatbot or agent internally. Employees could then ask real questions to test and genuinely understand the customer experience and drive improvements in the application of AI.”



Georgina Bulkeley

Director, Financial Services,
Global Strategic Industries, Google Cloud



The 2026 opportunity: A path to business growth





We have explored the five critical shifts in business driven by AI agents in 2026.

The redefined role of the employee as strategic orchestrator, the creation of “digital assembly lines” running core business processes, the rise of the agentic concierge, the move to agentic security operations, and the need for a human-centric upskilling strategy are changing the paths businesses are taking.

The curve for agentic systems might be steep for some organizations, but the true value is not just in the final product, but also in the innovation and optimization that happens during the process.

The companies that experiment today are not just building tools; they are building the critical, in-house expertise to manage, govern, and scale this new capability.

The 2026 opportunity can seem technical, but it is fundamentally human. It’s about freeing your teams from the repetitive, low-value work that drains their energy, allowing them to focus on the creative, strategic, and empathetic work that only they can do. This is your path to building a faster, smarter, and ultimately, more human company.

“ Access to agentic AI capabilities will democratize insights, innovation, creativity, and business growth, bringing value to consumers, employees, and organizations. This opportunity is truly significant, but comes with tremendous responsibility to ensure that the promise of AI delivers secure, ethical, and fair outcomes for all.”



Anil
Jain

Global Managing Director,
Global Strategic Industries,
Google Cloud

Google Cloud

See what
these trends
mean for your
business.

Get in touch

