Google Cloud Deloitte.

The automated enterprise:

Agentic AI and the new security imperative

“We’re entering a world where organizations may soon have more AI agents than human employees. This opens the door to new levels of scale, speed, and innovation — but it also creates massive complexity in governance and oversight. To truly take off with agentic AI, organizations must balance these opportunities and risks.”

Rayn Veerubhotla, Managing Director, Partner Engineering and Portfolio Strategy, Google Cloud

Introduction

What if your next hire wasn't a person? Imagine a digital workforce that never sleeps, executes complex strategies while you rest, and learns from each outcome. This is the world of agentic AI — a leap that has business leaders strapped in, engines roaring, as they prepare to taxi down the runway and soar into uncharted territory. But before the wheels lift off, the real challenge begins.

The business potential of agentic AI is vast, touching everything from HR to marketing, cybersecurity, and beyond. But first, leaders must ask themselves: Are we ready? Do we know what it takes to be ready? Agentic AI shows promise for unprecedented scale and innovation but also demands a new level of governance and security. Agentic AI is only as good as the controls around it, the systems beneath it, and the crew behind it. If your access policies are shaky and your teams aren't equipped to course-correct, AI agents won't fix problems — they'll create them.

Read on to learn what it takes to help ensure secure, strategic adoption of agentic AI — and how to stay in control at 10,000 feet.

Agentic AI: From automation to autonomy

Agentic AI represents the next quantum leap in the evolution of generative AI (gen AI), building on the capabilities that have recently captured the world's attention. Gen AI was designed to encapsulate human knowledge in a form that machines can understand, predicting what comes next — the next word, sentence, or idea — based on patterns learned from vast amounts of data. Agentic AI builds upon this foundation by taking those insights and autonomously making decisions, initiating actions, and performing tasks, with or without human oversight. Think of it as gen AI enhanced with mission logic, evolving from an assistant into a digital operator capable of dynamic problem-solving and autonomous action.

50%

of enterprises using gen AI are expected to deploy AI agents by 2027, up from 25% in 2025.[1]

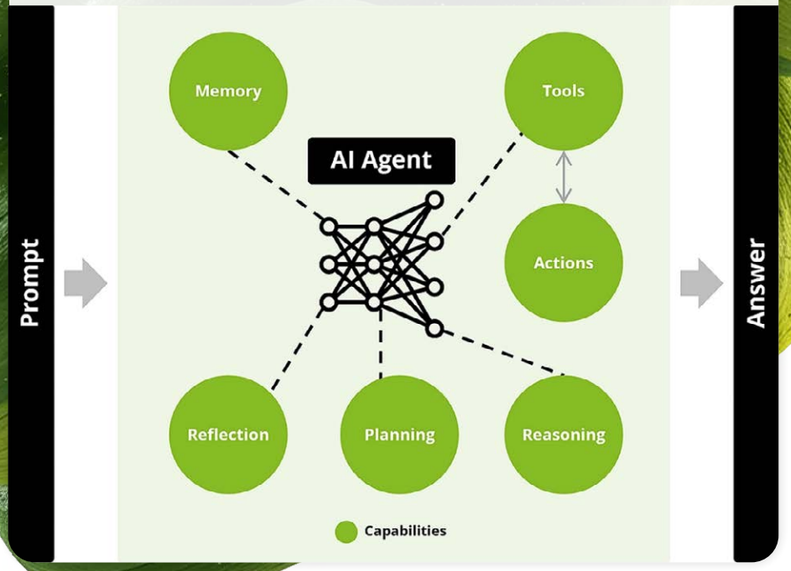

Image source: Deloitte CH, Agentic AI: The new frontier in AI evolution

1. Deloitte US, Deloitte 2025 Predictions Report, Generative AI: Paving the Way for a Transformative Future in Technology, Media and Telecommunications, Nov 17, 2024

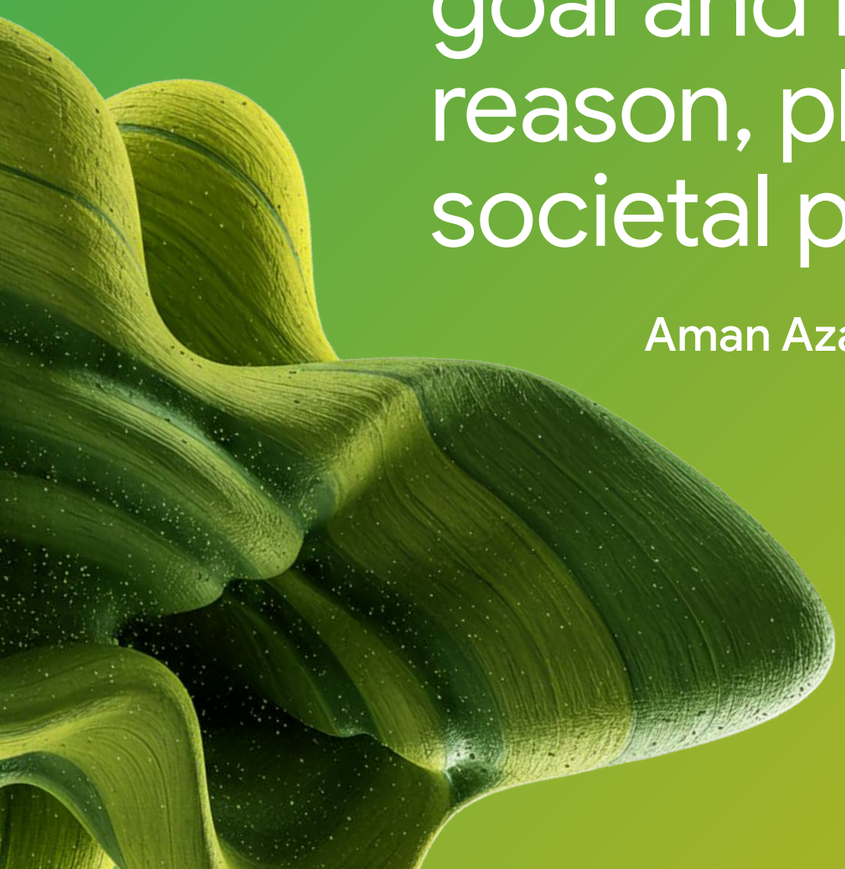

“Agentic AI can learn from outcomes, adapt, and improve over time. Simply give it a goal and it can perceive the environment, reason, plan, and act. The business and societal potential is immense.”

Aman Azad, AI Evangelist and Global Alliance Lead for AI, Google Cloud

Agentic AI in action

What does agentic AI look like in action? Given its goal-oriented nature and ability to plan and act, agentic AI can automate entire workflows across many business functions. Let's say an executive asks agentic AI to review upcoming marketing assets. The agent could automatically check them for brand compliance, flag inconsistencies, and suggest adjustments. AI agents can also streamline vendor evaluation, contract review, and approval across procurement processes. In finance, they could proactively manage working capital by monitoring cashflows, identifying bottlenecks, and recommending optimizations. In each case, it's like having a digital colleague who anticipates needs, ensures compliance, and keeps operations moving without being prompted.

As outlined in Google Cloud's *The transformative potential of agentic AI and the strategic imperative for Google Cloud partners* paper,[2] there are a variety of promising use cases, including:

- **Cybersecurity:** Help security teams quickly sort through hundreds of thousands of daily alerts to identify and triage cyberthreats and automatically recommend (or take) defensive actions.
- **Marketing:** Automate creating, personalizing, and optimizing multi-channel marketing campaigns for faster, more precise, and consistent results.
- **Returns:** Automate complex e-commerce return processes by coordinating warehouse, finance, and inventory systems without human intervention.
- **Inventory management:** Automate monitoring, restocking, and coordination across systems to optimize stock levels and prevent outages.
- **Customer experience:** Handle common questions and assist live agents by quickly gathering information, reducing call times, and improving customer satisfaction.

2. Google Cloud, The transformative potential of agentic AI and the strategic imperative for Google Cloud partners

“We are very pleased by these initial results, and we’re trending towards our goal of >95% automated verification. Document AI does what it says it can do, and we know we can leverage it economically to train for everything in the future.”

Kevin Cornish, Chief Information Officer, Covered California

COVERED CALIFORNIA

In 2023, Covered California ran a small pilot program with Google Cloud to test the viability of Document AI to meet its accuracy, efficiency, and scalability requirements. This generative AI-powered solution uses machine learning (ML) to automate the repetitive task of verifying resident information, improve the speed and accuracy of data extraction, and pull deeper insights from unstructured data contained in various documents. The team first used the solution out of the box without configuring it to establish a benchmark, achieving a document verification rate of 80-96% depending on document type for an average of 84%. Its legacy solution had an automated verification rate of 28-30%. With additional training, the new ML-based system’s success rates will continue to improve.

3. Google Cloud: *Covered California: Streamlining health insurance eligibility and verification with AI*

From takeoff to turbulence: Keeping agentic AI on course

There's no denying the power of agentic AI. However, with great power comes great responsibility. Agentic AI has the potential to deliver meaningful impact, but its effectiveness largely depends on the foundation on which it's built. For CISOs and CxOs, the question isn't just what AI agents can build, but whether the environment is ready to let them thrive. To realize the full potential of agentic AI, leaders should address a few obstacles first.

Turbulence ahead: Obstacles to building secure, scalable agentic AI

Once organizations address these risks, the benefits can be significant — and that's the reassuring side of the story. Engaging the right external collaborators with deep knowledge in both AI and cybersecurity can provide the guidance needed to build reliable agents and establish strong ROI. Without that groundwork, even the most advanced AI agents risk failure — with a potential for a costly rebuild just months after takeoff.

Access control and security:

Access control lists (ACLs) are typically buried deep within individual systems. But AI agents don't stay confined; they cut across systems and operate beyond traditional boundaries. How do you govern these agents? What permissions do they have? What guardrails are in place? These are highly complex matters, and governance and security are vital.

Skills gap and experience:

Organizations need very skilled people to build enterprise-grade agentic AI systems. But it's a new field, and there aren't enough skilled people out there yet. Agentic AI is a whole new ballgame, and many organizations are still warming up before they step up to the plate.

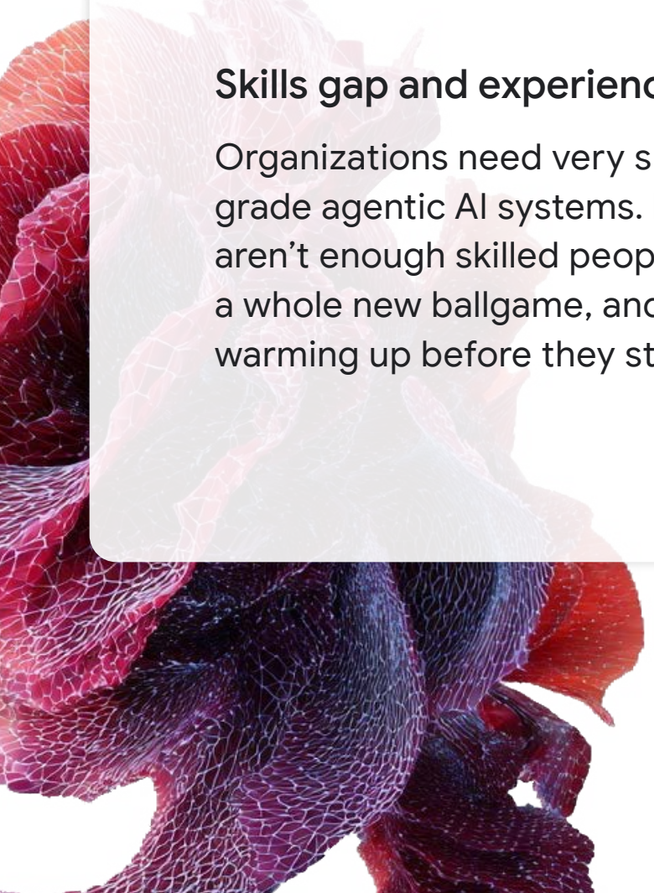

Hallucinations and cascading failures:

One challenge with gen AI is the risk of hallucinations, where the AI fabricates information based on the nearest approximation. AI can make mistakes because it's not an exact match — or an exact science — anymore. When agents "talk" to each other, one mistake, hallucination, or bad input can quickly snowball. Grounding models in enterprise data using technologies like Vertex AI Search can help ensure that responses are based on company-specific, factual information.

ROI and navigating the unknown:

While calculating the ROI of AI has been challenging, the outlook is improving. The cost is decreasing due to more efficient models, smaller architectures that perform comparably to larger ones, and advances like model distillation. Still, some fear the unknown. Many leaders question how autonomous agents will behave in complex, mission-critical scenarios.

Agentic AI is like an autopilot system that can rewrite its own flight plan. Every agent you deploy becomes a potential new attack vector, especially if they're given too much agency without proper governance. Without the right guardrails, it's easy to lose ground.

Gopal Srinivasan, Principal, Alphabet Google Alliance
Generative AI Leader, Deloitte Consulting LLP

Your flight plan: A framework for secure adoption

Every successful flight requires more than just engines and lift — it demands a flight plan, a skilled crew, and systems that anticipate turbulence before it hits. The same is true for adopting agentic AI. Deloitte and Google Cloud bring a structured methodology to help organizations chart a safe and scalable course.

At the center is Deloitte's Trustworthy AI™ Framework, which provides the governance and risk controls needed to keep AI initiatives aligned with enterprise strategy and regulatory expectations. This includes enterprise-wide access controls to help ensure only the right people — and the right agents — can access sensitive systems and data. Just as a pilot relies on accurate instrumentation, grounding AI in trusted enterprise data helps reduce bias and strengthen decision-making.

It's equally vital to introduce a human-in-the-loop review process at important checkpoints. This safeguard is essential: even as agentic AI operates at scale and speed, human oversight can ensure the system stays on course and mitigates risks before they escalate. Together, Deloitte and Google Cloud help enable organizations to adopt agentic AI with confidence, balancing innovation with control. With the right governance "flight plan," leaders can realize the transformative potential of AI while staying secure, compliant, and firmly in command.

2. Google Cloud, The transformative potential of agentic AI and the strategic imperative for Google Cloud partners

Innovation and control: Adopting agentic AI with confidence

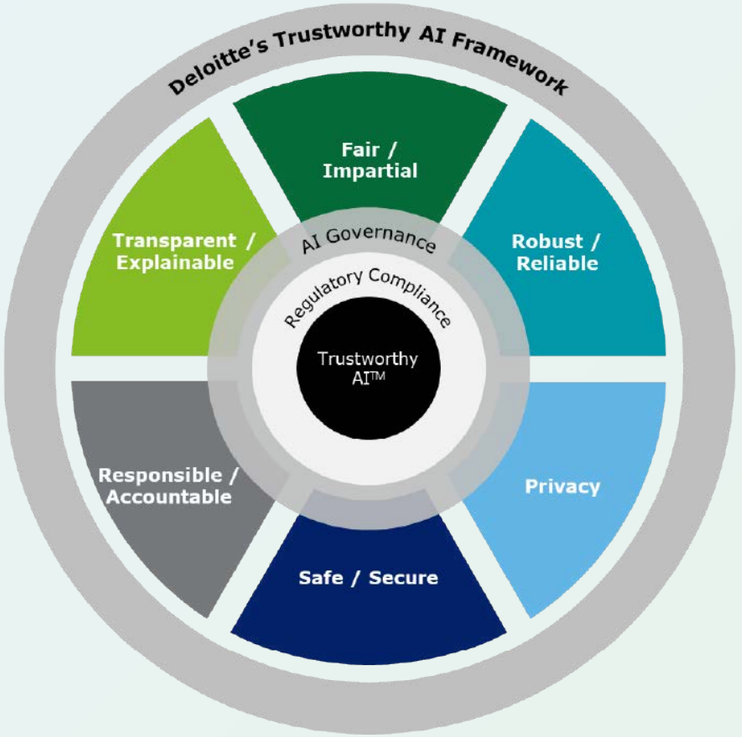

Image Source: Deloitte Development LLP, "Trustworthy AI™ Framework"

All systems go: Grounding agentic AI in reality

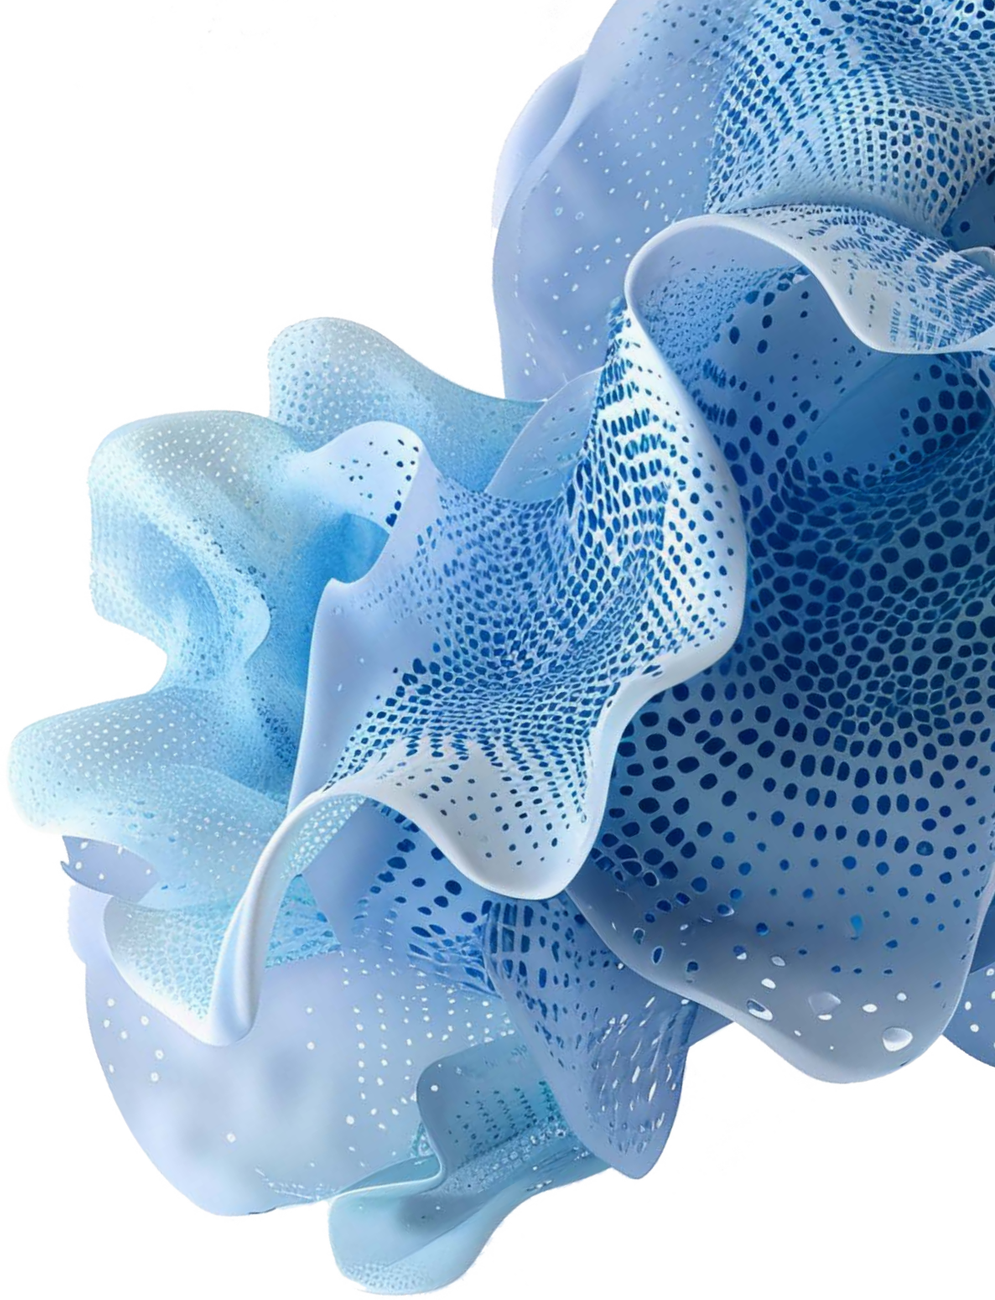

Agentic AI isn't just another tool in the enterprise tech stack; it's a new operating model for how work gets done. But success won't come from throwing caution to the wind, charging ahead recklessly without a plan. Nor will it come from sitting idle, waiting endlessly for perfect conditions to take off. This build-the-plane-as-we-fly-it moment can benefit from intentional piloting, a strong flight crew, and smart systems. The organizations that succeed won't necessarily be the fastest off the runway; they'll be the ones steadily gaining altitude through small wins, strengthening governance and security, and bringing the right crew along for the journey. When done right, the payoff isn't just operational efficiency or smarter automation — it's entering a whole new realm of innovation, where human and business potential are multiplied.

Google Cloud Deloitte.

Don’t let uncertainty ground your ambitions. An agentic enterprise requires a clear flight plan. Schedule your complimentary AI readiness assessment today to help ensure your strategy, security, and governance are cleared for takeoff.

Contact us

About

Google Cloud

Google Cloud is the new way to the cloud, providing AI, infrastructure, developer, data, security, and collaboration tools built for today and tomorrow. Google Cloud offers a powerful, fully integrated and optimized AI stack with its own planet-scale infrastructure, custom-built chips, generative AI models and development platform, as well as AI-powered applications, to help organizations transform. Customers in more than 200 countries and territories turn to Google Cloud as their trusted technology partner.

Deloitte

Deloitte provides leading professional services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets and enable clients to transform and thrive. Building on its 180-year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte's approximately 460,000 people worldwide make an impact that matters at www.deloitte.com.

Disclaimer

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (DTTL), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited (“DTTL”), its global network of member firms or their related entities (collectively, the “Deloitte organization”) is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication. DTTL and each of its member firms, and their related entities, are legally separate and independent entities.