

# Agentic AI & Generative AI: How Should They Function in Your AI Ecosystem?



Andy Maskin

Director, AI Creative Technology, Publicis Sapient



There's a major evolution happening in AI right now: generative AI to [agentic AI](#)—two different ways of making machines smarter and totally different. Generative AI is what *makes things*, like ChatGPT writing text or DALL-E creating images. But agentic AI? These digital assistants can actually *do things* on their own without needing constant instructions, almost like a co-worker or assistant.

While agentic AI offers greater potential power, it also brings complexity and integration challenges which slow down value creation and scalability. In contrast, generative AI's lower deployment barriers make it more immediately valuable, ensuring faster adoption in the global market for the foreseeable future.

How should businesses approach short-term and long-term investments in these two different categories of AI—when it comes to the internal enterprise and external, customer-facing use cases?

# Generative AI: A technical breakdown—and what the term actually means

## What is generative AI?

Generative AI refers to a class of machine learning (ML) models designed to produce new content—text, images, audio, code—by identifying and replicating patterns from extensive training data. These models, typically based on deep learning architectures like transformers or generative adversarial networks (GANs), generate outputs that align with the statistical properties of their training data rather than retrieving or modifying pre-existing content.

## Key use cases of generative AI

**Content creation:** Generative AI can produce new content—including text, imagery, audio and synthetic data—by learning from large datasets.

**Pattern recognition:** These systems identify and copy patterns in data, enabling them to generate outputs that resemble human-created content.

**Domain versatility:** Generative AI has applications across various domains, from art and entertainment to healthcare and finance.

## Automated content generation

One of the most prominent applications of generative AI is automated content generation. Tools like OpenAI's GPT-4o can generate human-like text, enabling applications such as:

**Writing assistance:** AI-powered writing assistants help users draft emails, articles and reports by providing suggestions or generating entire sections of text.

**Creative writing:** Authors and content creators use Generative AI to brainstorm ideas, develop storylines or even compose poetry.

## Generative AI for content creation

Beyond text, generative AI extends to other forms of content creation:

**Image generation:** AI models like DALL-E and Google's Imagen can create images from textual descriptions, allowing artists and designers to visualize concepts rapidly.

**Music composition:** AI systems can compose original music by learning from existing compositions.

## Agentic AI: A technical breakdown and what the term actually means

Agentic AI refers to systems designed to autonomously pursue complex goals with minimal human intervention. It's built on generative AI and applied in a different way. These AI agents exhibit autonomous decision-making, planning and adaptive execution to complete multi-step processes. Agentic AI can go into another computer system and solve the problem, unlike generative AI that can generate step-by-step instructions for you to solve a problem yourself.

### Key features of agentic AI

**Autonomy:** Agentic AI systems can make decisions and take actions without direct human oversight. They are capable of independently analyzing challenges, developing strategies and executing tasks.

**Adaptability:** These systems can adjust their behavior in response to changing environments and unforeseen events, relying on patterns and likelihoods to make decisions.

**Goal-oriented behavior:** Agentic AI is designed to accomplish specific objectives, often coordinating multiple sub-tasks as it communicates with other systems outside of itself.

### Examples of agentic AI applications

**Healthcare:** AI agents assist in registering participants for clinical trials, ensuring proper post-discharge care and aiding doctors in quickly understanding patients' medical histories. For instance, companies like [Grove AI](#) and [Infinitus Systems](#) have developed agents to handle administrative responsibilities, reducing physician burnout and managing workloads amidst doctor shortages.

**Customer service:** [ServiceNow has developed AI agents that handle 80 percent of customer interactions](#), with humans managing the more complex 20 percent. These agents have reduced handling times for complex cases by 52 percent and are projected to save \$325 million annually by increasing productivity.

## The main difference between generative AI and agentic AI

First, it is important to understand that agentic AI refers to an application of various AI technologies, **including generative AI**. It is not a specific AI technology, like generative AI. It is likely that most AI agents are based on generative AI (and other types of AI, like natural language processing) so that they can communicate with us in natural language, generate images or texts or help us brainstorm.

For an AI agent to be able to act on our behalf and make decisions, it will probably need to use a variety of technologies, like machine learning, natural language processing, systems integration, generative AI, deterministic AI and others to do things.

## **The market value of generative AI vs. agentic AI**

This means that agentic AI or AI agents will generally take longer to build, train and deploy successfully, because they are all so unique. For example, an AI agent that can make our doctors' appointments will function a lot differently than an AI agent that analyzes supply chain forecasting data to make real-time decisions to re-route goods.

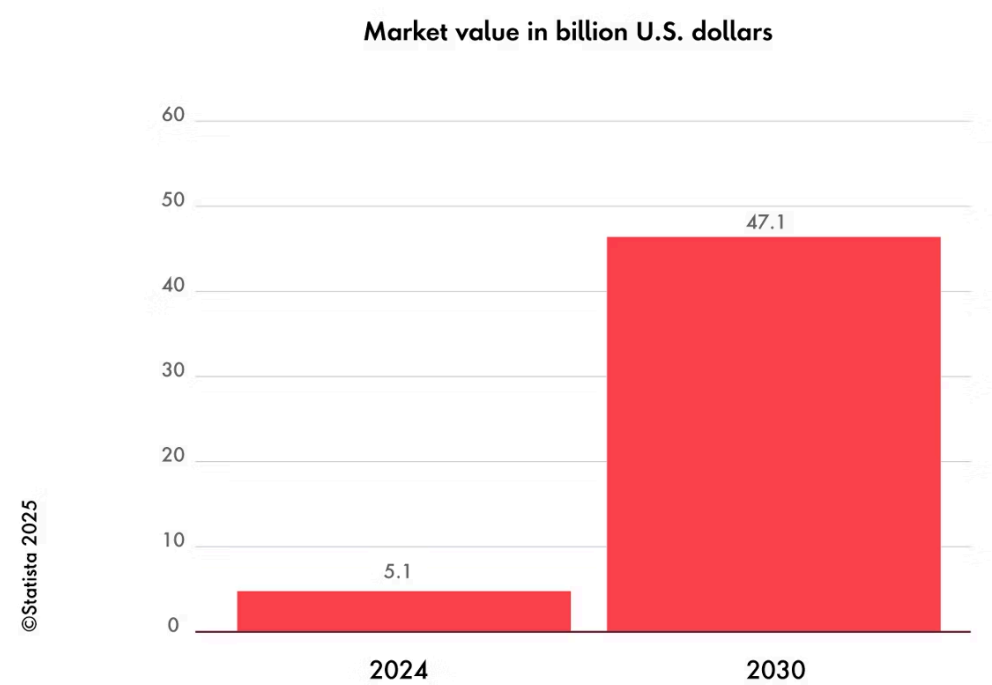
In each scenario, this AI agent or agentic platform will need to integrate with very different data sources and systems and have different guardrails and privacy rules set up to make decisions. While these requirements are also true of generative AI applications, the specific technological integrations are not necessarily required. Generative AI cannot take action, and therefore it doesn't need to integrate directly with platforms and accounts.

AI agents will theoretically have a higher output and impact on society because they are able to do the work for us. It also means it will be much more difficult to achieve and scale in the coming years.


Experts predict the global market value of agentic AI to [jump from \\$5.1 billion USD in 2024 to \\$47.1 billion USD in 2030 ↗](#), a 161 percent increase. In contrast, generative AI has an estimated [global market value of \\$36.06 billion USD in 2024, compared to \\$356.05 billion USD in 2030 ↗](#), a 163 percent increase. Generative AI has a much higher global market value

because it is faster and easier to scale for more generic applications, especially chatbots.

### The global market size of generative AI



### The global market size of agentic AI



# The future of AI: Implications of agentic AI and generative AI

Both agentic AI and generative AI continue to grow in market size at relatively the same rate, according to experts (161 percent vs. 163 percent). However, agentic AI applications will be limited to companies with flexible, composable technology architecture and higher data maturity. Similarly, widespread acceptance of agentic AI will take more time to develop.

If we compare the acceptance of self-driving cars to agentic AI, the pattern is similar—despite advances in technology, most people still don't fully trust them. At the same time, their success depends on working seamlessly with other systems, from public transportation and ride-sharing services to partnerships with car manufacturers.

Because of this phenomenon, many businesses are choosing to invest in hyper-specific, internal use cases for agentic AI, that don't rely on external integration or customer buy-in.

## Agentic AI

For example, in healthcare, AI agents can manage administrative tasks, allowing medical professionals to focus on patient care, by integrating natural language processing (NLP), machine learning (ML), robotic process automation (RPA) and rule-based decision engines.

The AI agent would then integrate with hospital electronic health records (EHRs) using **GPT-like transformer models** to create concise, structured summaries for progress notes, discharge summaries and referrals.

These are all of the technical challenges and solutions that this use case would require:

## Challenge

- Unstructured medical data
- EHR interoperability
- Regulatory compliance
- Data privacy
- Prescription fraud prevention

## AI solution

NLP-based entity recognition

FHIR/HL7 integration

AI-driven policy rule enforcement

Federated learning for decentralized model training

Blockchain-based audit trails

## Generative AI

In contrast, a generative AI-only healthcare solution would require more human intervention and medical expertise, and may save less time on behalf of a medical professional, but would be faster to implement.

A purely generative AI solution would be faster to implement because it relies on pre-trained language models that can be fine-tuned on medical data, rather than requiring complex rule-based systems and deterministic AI pipelines.

It eliminates the need for structured data integration, hardcoded compliance checks and real-time cross-referencing with external medical databases, reducing development time. Additionally, generative AI can be deployed as a conversational interface with minimal backend changes, while a fully agentic solution requires deep integration with EHRs, pharmacy networks and regulatory systems.

## When to invest in generative AI solutions vs. agentic AI solutions

Developing a fully independent AI agent or [workflow of AI agents](#) will be more complex and time consuming than developing a generative AI solution for a variety of reasons, from technology integrations to data privacy, to deterministic or hardcoded elements and more.

For especially complex, time-intensive and costly workflows that:

- a) Rely on analyzing vast amounts of data very quickly
- b) Are absolutely essential to your business model

c) Need to happen in real-time very quickly,  
investing in a proprietary AI agent will bring significant value.

## When to use third-party agent solutions

However, most companies won't have the money or time to invest in their own proprietary AI agent, at least not right now. Third-party tools can be a practical alternative for automating non-core tasks. These platforms offer ready-made AI agents (like [Agentforce](#)) which can quickly handle customer service chats, document processing or internal knowledge management with only minor customization.

While they may lack deep system integration, they still provide efficiency gains without long development cycles. For critical functions, a custom-built agent might be worth the investment, but for standardized, repeatable workflows, a third-party tool could get the job done faster and at a lower cost.

## An agentic solution: Sapient Slingshot

Publicis Sapient's AI product [Sapient Slingshot](#) is an example of when the creation of a proprietary AI agent (actually an ecosystem of AI agents) was worth the investment. This platform accelerates enterprise system integration and software development by using AI agents to automate code generation, testing and deployment, reducing project timelines from months to weeks.

### Why was this necessary?

Sapient Slingshot draws on 20+ years of internal code developed by thousands of employees to quickly process millions of lines of code. Software development, application development and legacy application modernization, all of which follow [the software development lifecycle \(SDLC\)](#), are essential to our business model and the service and value we provide to our clients. These projects typically take many years, and we can reduce the timing in half with this solution.

**Generative AI alone was not effective for Sapient Slingshot** because system integration requires precise execution of APIs, data transformations and



compliance with enterprise IT architectures, which LLMs cannot reliably enforce. Sapient Slingshot needs structured automation for tasks like code generation, testing and deployment, ensuring accuracy, security and performance—something generative AI struggles with due to its probabilistic nature and inability to validate outputs against system constraints.

**A third-party AI development tool also wasn't a fit** because many off-the-shelf code assistant solutions lacked the customization, security and integration needed for enterprise-scale system orchestration. Legacy modernization and software development are core to Publicis Sapient's business, and our approach is uniquely tailored to complex enterprise environments. Third-party tools can't adapt to our proprietary workflows or handle the precision required for code generation, testing and deployment. Building Slingshot in-house ensures we maintain full control, optimize for our specific needs and deliver the reliability and scalability that generic solutions can't provide.

## Top agentic and generative AI use cases across industries




While we are still in the early stages of agentic AI exploration and development, we are already seeing use cases across industries where there are significant opportunities for companies to begin developing AI agents, or utilizing third-party applications of AI agents that can integrate with proprietary data.

**Here are our top ideas:**

### Agentic AI use cases across industries


## Retail

**Use case:** An AI agent that monitors real-time sales data, supply chain disruptions and local consumer demand to automatically adjust pricing and restocking schedules for individual store locations.

**Why it's worth it:** [Retailers lose \\$1 trillion annually due to stockouts and overstocking](#) . A dynamic AI pricing and inventory agent could optimize margins, reduce waste and boost revenue by 5-15 percent per retailer—leading to billions in new profit across the industry.

## Consumer products

**Use case:** An AI agent that predicts shifts in consumer demand for fast-moving consumer goods (FMCG) based on social media trends, weather data and economic factors, automatically adjusting production schedules and supplier orders.

**Why it's worth it:** [Consumer brands face \\$250 billion in lost revenue annually due to inaccurate demand forecasting](#) . This AI agent would reduce overproduction costs, prevent stockouts and increase revenue through more accurate supply chain adjustments, making it an essential investment for manufacturers and distributors.

## Financial services

**Use case:** A real-time financial assistant that monitors customers' spending, predicts cash flow issues and proactively recommends personalized loan options while auto-filling applications and checking risk factors.

**Why it's worth it:** Banks struggle with high default rates and low customer engagement. This AI agent would increase loan approval efficiency, drive revenue through personalized lending offers and reduce risk—potentially adding billions in new lending opportunities while lowering operational costs.

## Energy and commodities

**Use case:** An AI agent that automates carbon credit trading by analyzing real-time emissions data from factories, verifying compliance reports, and recommending optimal credit purchases or sales.

**Why it's worth it:** [The carbon credit market is expected to reach \\$100 billion by 2030](#) . Automating compliance and trading with AI would unlock

untapped liquidity, prevent regulatory fines and create a new revenue stream for energy companies—making this an essential profit-generating tool.


## Travel and dining

**Use case:** An AI agent that predicts customer foot traffic in restaurants, hotels and airlines based on weather, local events, online bookings and historical data, then adjusts labor schedules and supply orders automatically.

**Why it's worth it:** The hospitality industry loses billions annually due to labor inefficiencies, food waste and staff shortages. AI-driven scheduling could cut labor costs by 20 percent, increase staff efficiency and boost revenue through optimized staffing—leading to billions in cost savings and improved guest experiences.

## Transportation and mobility

**Use case:** An AI agent that monitors vehicle sensor data, predicts mechanical failures before they happen and automatically optimizes delivery routes based on real-time traffic, weather and fuel efficiency models.

**Why it's worth it:** Transportation companies lose [\\$50 billion annually due to unexpected vehicle downtime and inefficient routing](#) . This AI agent could cut maintenance costs by roughly 30 percent and fuel costs by 10-15 percent, leading to billions in industry-wide savings and increased delivery efficiency.

## Public sector

**Use case:** An AI agent that automates fraud detection and real-time claims verification for social security, unemployment benefits and tax refunds by cross-referencing applicant data with banking transactions, employment history and identity verification databases.

**Why it's worth it:** Governments lose billions annually to fraudulent claims, improper payments and administrative inefficiencies. This AI agent would reduce fraud, speed up benefit approvals and lower operational costs by automating document verification and flagging anomalies—potentially saving tens of billions globally each year.

## Health

**Use case:** An AI agent that automates prior authorization approvals for insurance claims by reading patient history, validating medical necessity and auto-filling forms for providers and payers.

**Why it's worth it:** Prior authorizations delay treatments, cost insurers and providers billions annually and frustrate patients. Automating this process could reduce administration costs by 50 percent, improve patient outcomes and create a multi-billion-dollar AI-driven insurance automation market.

## **Generative AI use cases across industries**

At the same time, there are valuable opportunities for generative AI applications across industries that can provide immense value on a faster timeline.

This is not to say that there is no risk or complexity involved, or that these solutions can be built overnight. These are just products where we are already seeing success with clients that have been able to generate ROI and customers and employees alike are gaining value and time from using them.

### **Retail**

**Use case:** A generative AI model that automatically writes product descriptions based on key features and summarizes thousands of customer reviews into easy-to-read insights.

**Why it's worth it:** Requires no enterprise system changes, just plugging into existing e-commerce platforms (e.g., Shopify, Amazon). It boosts conversion rates and reduces manual content creation costs immediately.

### **Consumer products**

**Use case:** A generative AI tool that writes product marketing copy, personalized email campaigns and social media content based on customer behavior and preferences.

**Why it's worth it:** No data security concerns since it works with existing CRM and marketing automation tools (e.g., Salesforce, HubSpot). It reduces marketing costs while improving engagement rates instantly.

### **Financial services**

**Use case:** A generative AI system that auto-generates responses to customer inquiries, summarizes financial reports and creates customer-friendly explanations of complex financial policies.

**Why it's worth it:** No financial decision-making or compliance processing—just improving clarity and efficiency. Banks and insurers can reduce customer support costs while enhancing customer trust.

## Energy and commodities

**Use case:** A generative AI assistant that automates Environmental, Social and Governance (ESG) reports, summarizes regulatory compliance changes and generates investor-friendly sustainability disclosures.

**Why it's worth it:** It doesn't alter operations, just speeds up reporting and reduces legal review overhead—helping companies meet ESG demands faster and improve investor transparency.

## Travel and dining

**Use case:** A generative AI tool that creates personalized travel itineraries for customers and automatically generates responses to online reviews, improving brand reputation.

**Why it's worth it:** No customer data processing—it just repackages public travel content and assists hospitality teams with high-volume review management, increasing bookings and loyalty.

## Transportation and mobility

**Use case:** A generative AI-powered logistics email assistant that drafts responses to shipment inquiries, delays and tracking updates for logistics and freight companies.

**Why it's worth it:** No backend system changes, just auto-generated, pre-reviewed customer communications that reduce response time and increase satisfaction with zero security risks.

## Public sector

**Use case:** A generative AI chatbot that answers common citizen inquiries (e.g., tax filings, benefit applications, permit requests) and generates personalized responses based on official government data.

**Why it's worth it:** No sensitive data processing—just publicly available information restructured for clarity. It reduces call center workload and improves response time for high-volume inquiries.

## Health

**Use case:** A generative AI-powered medical scribe that transcribes patient visits in real time, automatically summarizing key points into structured notes for EHR entry.

**Why it's worth it:** No decision-making or diagnosis involved, just automating documentation. Physicians save hours per week on paperwork, improving patient interaction time without liability concerns.

## Keeping the "human in -the loop" with agentic and generative AI

Contrary to the "independent" nature of agentic AI, both generative and agentic AI (but especially agentic AI) require a "human in -the loop." Human intervention is essential in the model development, training, usage and review.

Just because AI can be better and faster than humans with certain tasks or use cases, when things go wrong (a generative AI model hallucinates, or an AI agent makes a bad decision), us humans are to blame, not the technology.

When we as people fully understand our AI solutions—including the benefits and risks, as well as plans and policies in place for a variety of unexpected situations—only then will we be able to gain the true value from them. Before we build, we should ask ourselves, our employees and our customers: What are the pros and cons of this solution? How will this help you? What are your thoughts? What are the risks? What do you need?

If we were to put a child behind the wheel of a car, and they hit a tree, we would not arrest the car. We would talk to the child, and his or her parents. But that does not mean we go so far as to say "No one should ever drive a car." The same is true of AI technology. If an AI tool tells a person to file their taxes incorrectly and they face a penalty, it is up to the business that created the tool to fairly compensate this person, take responsibility and fix the situation quickly.

As these technologies evolve, mistakes in their implementation and use are inevitable. However, keeping humans in the loop at every stage helps mitigate negative consequences and allows us to adapt and improve. At the same

time, choosing to ignore agentic and generative AI carries its own risks—just with different trade-offs.

## Related Topics

---

[Artificial Intelligence](#) | [Machine Learning](#)

© 2025 Publicis Sapient. All rights reserved. A Publicis Groupe Company.