ORACLE

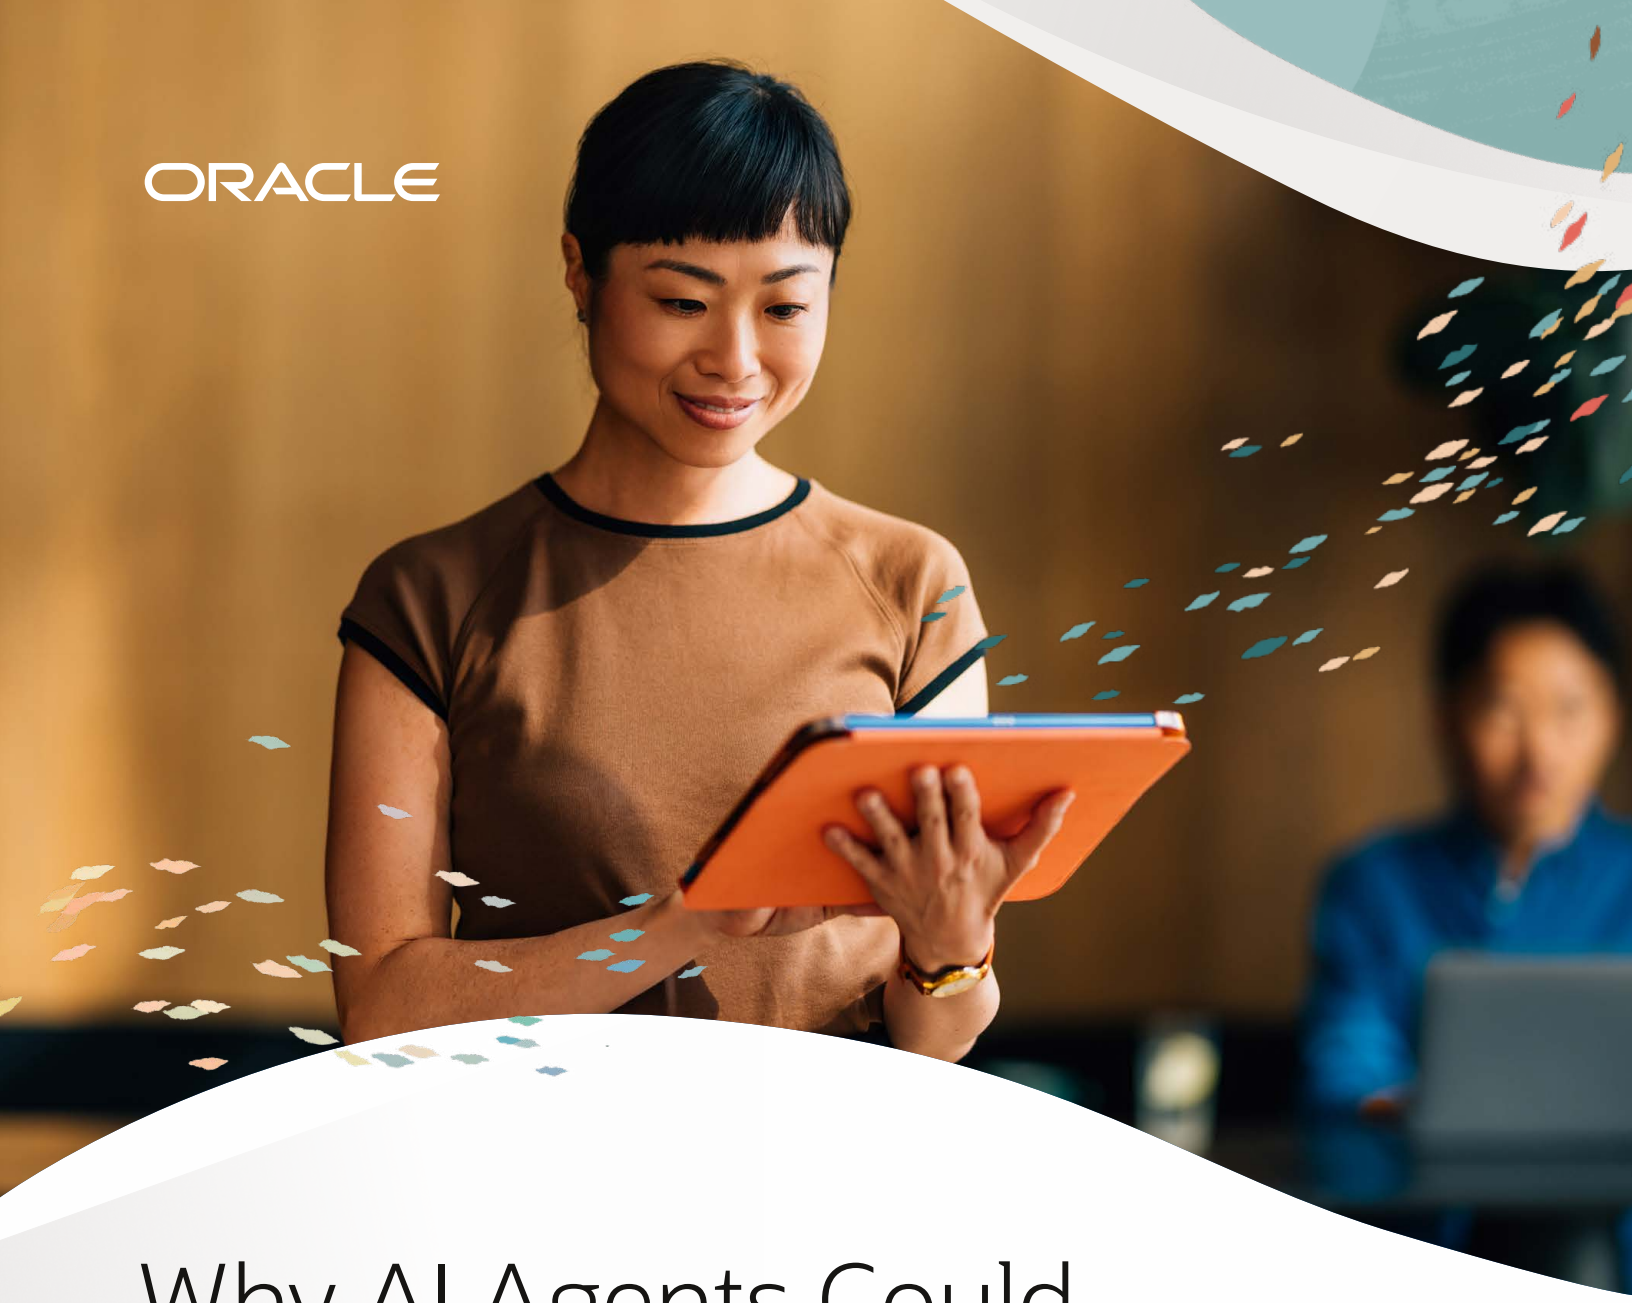

Why AI Agents Could Usher In the Next Productivity Boom

Software tools built on the latest LLMs and embedded in business applications are starting to deliver long-promised automation.

Table of Contents

By Aaron Ricadela
Senior Writer

In 1987, PCs loaded with spreadsheets, word processors, desktop publishing programs, and databases had permeated offices by the millions, replacing paper ledgers, filing drawers, and layouts made with scissors and glue. Yet US productivity growth was still inching along from early '70s levels, prompting economist Robert Solow, who won a Nobel Prize that year for proving a link between technology advancements and economic growth, to quip that the computer age was everywhere but in the productivity statistics.

By the mid-1990s, though, GDP growth per hours worked had doubled, as the internet and computer hardware advancements helped to transform office work and most every other business function. The lesson? It's always taken time for the effects of new technologies to show up in productivity numbers as they reduce the time it takes workers to complete some tasks, create space for new ones, and gradually improve processes.

Now, amid productivity growth that's slipped back to '80s and early '90s levels, generative AI promises similar improvements, even as executives and investors question whether this year's planned $320 billion in artificial intelligence–related data center spending will translate into meaningful business benefits.

As in decades past, an AI productivity boom may not happen overnight. One development that could spark it are emerging AI-powered "agents." These nearly self-driving software tools are designed to respond to conversational cues instead of code, automatically pull data

from a range of sources, and execute actions in business applications and on the web to help complete accounting, HR, sales, supply chain, and other administrative tasks.

While AI agents are still in their early days of deployment, 67% of top executives at companies worldwide consider them part of their AI transformations, with optimism "consistent across geographies," according to a Boston Consulting Group survey of 1,803 C-level execs in 19 countries.[1] Investment bank Barclays estimates agents could automate about 7 billion business tasks—2.5% of them in 2025 and double that percentage in 2026—helping teams and entire companies become more productive.[2]

> "Low- and no-code development tools are steaming ahead right now. They're the right place to build AI agents—business users are already familiar with them from other coding work."
>
> **Holger Mueller**
> Analyst, Constellation Research

The applications are myriad. Ford is reportedly using AI agents to turn 2D design sketches into 3D car models without having to build every prototype in clay. Home electronics company Sonos deploys agents that recall details of customers' conversations to help them set up its products, including troubleshooting home networking gear the company doesn't make. Healthcare group AtlantiCare is using AI agents to help doctors navigate electronic health records and extract information from clinical notes.

AI agents are poised to help customers reimagine all their front- and back-office applications, becoming a new business logic layer that executes transactions and queries. On the front end, they're augmenting point-and-click or command-driven user interfaces with the natural, typed, or spoken back-and-forth parlance computer users employ to converse with AI chat programs. Unlike previous generations of software assistants and automation tools, agents call on large language models (LLMs) to predict users' intent and complete code and forms, choosing the best tools to execute the job. That can make them less error prone and less reliant on hard coding.

But agents won't just hit the ground running. They're generally customized by organizations and need specific instructions and configuration to work. Since out-of-the-box LLMs don't

contain the company-specific data that makes agents shine, software vendors are supplying toolkits to design workflows and prepare documents that augment what AI models learned during their initial training. Organizations that deploy agents from different software suppliers may find themselves with an array of design studios.

What's more, industry standards for how agents from different vendors communicate are just now emerging; it's early days of adoption. And companies can't just assume their data is ready for AI. They need to make sure it's clean and stored in the right place.

Writing in *Harvard Business Review*, management guru Tom Davenport and data quality consultant Thomas Redman argue that while manufacturing process improvements, such as lean production and Six Sigma, were widely adopted in the 1990s and saved organizations billions, they fell from favor amid failed projects, overdone job cuts, and lack of IT support. AI offers companies a second chance to scale up process reengineering by helping find inefficiencies faster and creating process templates that suggest to managers how to optimize steps and communicate better with IT.[3]

To enter the AI agent era, companies will need to unleash them on processes that are already healthy, tune up those that aren't, and make sure their LLMs and agents connect to business applications and databases. On the people side, it's important for organizations to make sure that less skilled employees, and not just high performers, get value from these tools.[4] And boards of directors may consider what regulators expect from autonomous software that touches critical processes.

Read on to learn about AI agent use cases, options for deploying them, the importance of data preparation, and other key considerations.

1. Agents everywhere

Unlike LLMs that underlie AI-driven chat services, agents mediate between enterprise software users and underlying AI models. They can supplement LLMs' innate knowledge with a businesses' own information, adapt to changing circumstances, set priorities, and carry out time-saving steps.

Agents can pull data from finance, HR, sales, supply chain, and other enterprise applications, as well as from emails and calendars, showing computer users insights from across multiple systems. They can also tap structured information in databases and unstructured data from documents via retrieval-augmented generation (RAG), which supplies an LLM with access to proprietary business data as it's serving up answers.

Set up properly, agents even know when to pause and ask users for more information needed to complete a job. "Adept at both code and natural language, agentic systems will let people shift from operating software themselves to making requests that kickstart an agentic process that reasons," consultancy Accenture said in its Technology Vision 2025 report.[5]

The jobs agentic AI can tackle generally fall into three categories: Fast to deploy, able to answer questions, and capable of taking action.

To notch quick wins, an HR department might set up agents to create job postings or set performance goals. A marketing group could build them to create personalized customer emails and online landing pages. A supply chain team could use agents to generate shift summaries or instructions for inspecting goods.

The next level of sophistication is agents that can help answer questions—say, on employee benefits and pay slips, products and accounts, or material-handling guidelines. Agents that take actions can assist with onboarding new hires and help managers make pay decisions.

They can also triage service requests and plan steps needed to resolve them or process price changes and returns.

Other use cases lie in credit underwriting or showing potential financial results versus risk tolerance. Agents can also ease insurance paperwork for doctors and patients, analyze damaged equipment via photos, and research customers for sales.

Most enterprise software will be enhanced by agents in 2028 and 2029, market researcher IDC said in a February 2025 report.[6] The goal is for them to handle entire functions, such as supply chain management, logistics, and inventory management, replacing conventional user interfaces. AI assistants and advisers "have both quickly become must-haves in modern software," IDC said.

In addition to offerings from enterprise application vendors, including Oracle, SAP, Salesforce, ServiceNow, and Workday, Microsoft is building AI agents into Word, Excel,

Most Common GenAI Use Cases

Percentage of individuals in each corporate function reporting use.

16% **Content marketing support**

15% **Personalized marketing**

10% **Design development**

7% **IT help desk chatbot**

Source: *Feb.-March 2024 survey of 1,363 employees* by McKinsey QuantumBlack

Outlook, and Windows. Apple is working on an update to its Siri voice assistant that uses generative AI to carry out tasks such as filling out forms and planning an appointment.

Some of these vendors are outfitting customers with design studios for customizing, instructing, and activating AI agents—or creating their own. Oracle AI Agent Studio, for example, works with the company's Fusion Cloud Applications and APIs to help create agents from preloaded libraries or craft new ones. "The business user is building the agent," says Holger Mueller, an analyst at Constellation Research. "Low- and no-code development tools are steaming ahead right now. They're the right place to build AI agents—business users are already familiar with them from other coding work."

Well-funded AI startups are also vying for a piece of the market. OpenAI in March 2025 released APIs and tools to make building agent-driven applications simpler, with less prompt engineering and custom business logic required. Anthropic is working on features that would let its Claude models help knowledge workers with tasks in spreadsheets and documents, and it's published extensive guidance for businesses building agents.

2. Build vs. buy

One of the first decisions IT organizations need to make when considering AI agents is whether they want to customize the ones prebuilt by their companies' software vendors or build their own from scratch. Most will likely tailor agents from templates to start harvesting returns sooner and minimize the number of software components they need to build and manage.

The good news is that common use cases for automating back-office processes or carrying on conversations with customers don't necessarily need a lot of technical expertise to implement. Application administrators working with design tools can customize out-of-the-box agents, choosing from a select number of preintegrated LLMs. Since agents' performance and the syntax users need to prompt them are closely tied to underlying AI models, businesses should look for software-as-a-service providers that continually test how their agent environments perform and incorporate model changes in new releases of their applications.

Companies customizing prebuilt models can choose from templates in a catalog that contain the necessary code to run them, then type natural language instructions into a field or pick actions from lists to show the agents they're creating how to converse with users, display data, or set up appointments. Admins can also select which tools and documents agents should refer to when they're carrying out tasks or conveying information. To make sure that the agents are working as expected, AI agent studios contain simple testing tools to run through sample interactions and judge the results before deploying the software live.

Building new agents from scratch gives businesses more leeway to define specialized processes. But it requires more expertise in programming, sourcing and preparing data, and training users. Many businesses also lack the knowledge to select an LLM and fine-tune its performance. Commercial and open source software vendors offer AI agent frameworks, such as LangChain, LlamaIndex, and AutoGen, which handle low-level work, such as calling LLMs and functions in databases, calendars, or email programs.

Pricing for deploying AI agents also varies widely. Some ERP, HR, and CRM application vendors offer premium subscriptions to unlock more AI use cases, while others charge customers based on the number of transactions their agents complete. Businesses can look for vendors that include the functions in their cloud applications at no extra cost. Those building their own agents from the ground up need to account for the cost of API calls to LLMs and cloud infrastructure charges.

3. AI-ready data

AI agents can communicate among systems and retain useful information in memory to guide employees and customers and help complete tasks. To reap those advantages, data from ERP, HR, sales, supply chain, and other systems should retain descriptive metadata once the information is exported to data warehouses or lakehouses, before it's made available for AI. For example, descriptions can get lost when companies move information created in enterprise applications to external stores, making day-to-day tasks harder for departments to complete.

> "If it's not obvious to a bystander what the data means—which it isn't in most cases—then agents won't help you."
>
> **Miranda Nash**
> Oracle group vice president for applications development and strategy

Keeping application data in place helps ensure that IT departments can manage access by employees' roles. And context-enhancing RAG technology benefits from clean, unified information. Furthermore, the cost of cleansing and unifying data strewn across different systems and formats can run into the millions of dollars.

"If we had self-describing data everywhere, then we wouldn't have this issue," says Miranda Nash, Oracle group vice president for applications development and strategy. "But that's not the reality. If it's not obvious to a bystander what the data means—which it isn't in most cases—then agents won't help you."

Process management is also critical to success. To make the marriage of AI agents and operations reengineering work, businesses need to apply the technology to well-oiled processes that are already integrated end to end, which lets them harvest large amounts of

high-quality data needed for good outputs, Davenport and Redman write in their HBR article. “Good process management requires departments to adopt common data standards and share data freely across an organization, although many departments may be reluctant to do so,” the authors say.

IT architectures using agents to span processes in different applications can be changed more readily than ones that use so-called robotic process automation software. Because RPA links various applications according to hard-coded rules, businesses often need to change systems all at once, a situation IT departments generally try to avoid.

Agents are also changing the way companies think about getting a single right answer from ERP applications, says Kristian Kersting, a professor of AI and machine learning at Germany’s Technical University of Darmstadt and head of the Hessian Center for Artificial Intelligence. Unlike with conventional SQL database queries, prompts to agents don’t come with a guarantee of one correct return. For example, if an API call to an enterprise system used in RAG incorporates technical terms in its documentation that differ from the language a marketing department or user interface employed for the same concept, the LLM can return mistakes or unexpected behavior.[7]

“RAG says, ‘Produce your answer chiefly with this content,’” Kersting says. “That has advantages and disadvantages. The content has to be correct. But the UI is—and I don’t like this term—more ‘human’.”

4. From investment to productivity

Generative AI looks like what economists call a general-purpose technology: pervasive across sectors, rapidly improving, and spawning related innovations. Taking advantage of such breakthroughs (think: steam power, electrification, the microprocessor, and the world wide web) historically takes time, as companies rebuild processes and rebalance workforces.

Between 20% and 40% of US workers are employing AI on the job, according to a US Federal Reserve Board roundup of surveys released in February 2025.[8] In Europe it's more like 13.5%.[9] Yet just 1% of C-suite respondents described their GenAI efforts as mature, and most organizations aren't getting the ROI they'd hoped for, according to a McKinsey & Co. survey of 3,851 US C-level executives, managers, and other staff.[10]

One way to accelerate returns may be to prioritize the places where AI agents run. Companies that focus their AI efforts on an average of 3.5 pilot projects realize twice the ROI of those that spread their bets across at least six, according to Boston Consulting's global C-suite survey.

Ways AI Value Is Tracked

28% **Operational value only**

16% **Financial value only**

24% **Both operational and financial value**

32% **Not yet tracking**

Source: *2025 survey of 1,803 C-level executives worldwide* by Boston Consulting Group

And while research has shown that relatively inexperienced workers can get the biggest boost from GenAI, only 29% of survey respondents in 19 countries said they'd trained more than a quarter of their workforce on the tools, though the percentages were much higher in Singapore (44%) and Japan (38%).

Rapid computing advancements are also paving the way for wider agent adoption. A new generation of so-called reasoning models is better at breaking requests into discrete steps. New generations of AI chips, such as GPUs based on NVIDIA's upcoming Vera Rubin design, are designed to deliver better inference performance that agents need to do their jobs. LLMs' context windows, or the number of basic processing unit tokens they can handle, are growing quickly as well.

On the cost side, there's a movement toward smaller, more efficient language models that use less computing power and excel at specialized tasks. If business applications could combine larger, general-purpose models with smaller, specialized ones on customers' sites, it could improve performance with only modest increases in computing.

5. Nascent standards

Development of AI agents is moving so quickly that the tech industry hasn't yet established standards for managing data access, regulatory compliance, and security across different development frameworks, nor are there standards for how agents communicate with one another. Teams within the same business often adopt different frameworks and tools, making AI agent strategies difficult to scale. Portability across cloud platforms is still manual and error prone.

Tech companies are starting to work on solutions. An open source communications protocol called the Model Context Protocol (MCP), developed by Anthropic, provides a standard way for AI models to pull data from business software, helping speed development of agent systems. It's an alternative to businesses needing to write custom integrations to connect LLMs to each of their data sources. OpenAI is supporting the protocol in its SDK for agents and other products, and Microsoft is supporting it in its Copilot Studio.

Oracle is working on a standard format called Agent Intermediate Representation for creating portable AI agents that interoperate across different frameworks. IBM Research's Agent Communication Protocol (ACP) proposes standards for how open source agents communicate with one another.

Coming Capabilities

Percentage of organizations that expect AI agents to perform these functions in the next three years.

45% **Access internal IT functions**

45% **Automate tasks**

33% **Build new digital functions**

32% **Automate cross-company workflows**

29% **Access functions from other systems**

Source: *2025 survey of 4,021 global C- and director-level executives* by Accenture

6. Buttoning down governance

As agents get embedded more deeply into companies' operations, they're drawing more attention from executives and boards. Financial regulators in the UK, Singapore, and Hong Kong have said company directors need to oversee AI-related risks, such as data protection and human controls on models' output.

In the US, corporate case law related to the Caremark standard, which holds corporate directors accountable if they breach their oversight duties, has put AI use higher on boards' agendas. Courts have already examined claims against boards for insufficient cybersecurity oversight, and AI could be next, according to a 2024 analysis by financial communications advisor Edelman Smithfield.[11]

Businesses are still putting their oversight apparatus in place. At 28% of companies using AI, the CEO is responsible for overseeing governance of the technology, according to a July 2024 online survey of 1,491 employees of companies in more than 100 countries by McKinsey QuantumBlack.[12] Just 17% of respondents said their boards oversee AI governance.

Market researcher Evident reported in September 2024 that 41 of the 50 banks on its AI adoption and maturity leaderboard have an AI leader on their executive committee, compared with 38 in 2023. But most banks are spreading those responsibilities, Evident said, rather than creating new positions expressly to govern AI risk.[13]

How Oracle can help

Oracle AI Agent Studio lets organizations set up and deploy AI agents in their Oracle Cloud Applications, assigning them tasks, data access rights, and guidelines for interacting with users, and then testing the results. Designers can create agents from templates in a library or design entirely new ones, including the ability to tap OpenAI's GPT-4 hosted on Microsoft Azure.

Oracle Fusion Cloud Enterprise Resource Planning contains agents to help complete accounts payable, general ledger, reconciliation, and other finance processes.

Oracle Fusion Cloud Human Capital Management contains AI agents that respond to questions on employee pay and benefits, set performance goals, draft job, and onboard new hires.

Oracle Fusion Cloud Supply Chain Management & Manufacturing agents include those for creating inspection instructions for goods, drawing up guidelines for handling materials, and analyzing photos to estimate and start repairs on machinery and equipment.

Oracle Fusion Cloud Customer Experience includes agents that can research documents to see where sales accounts stand, identify upsell and renewal opportunities, and suggest customer discounts.

Oracle Cloud Infrastructure Generative AI Agents is a cloud service that lets businesses build LLM- and RAG-powered agents that can search their data stores to answer users' questions while documenting the information's source.

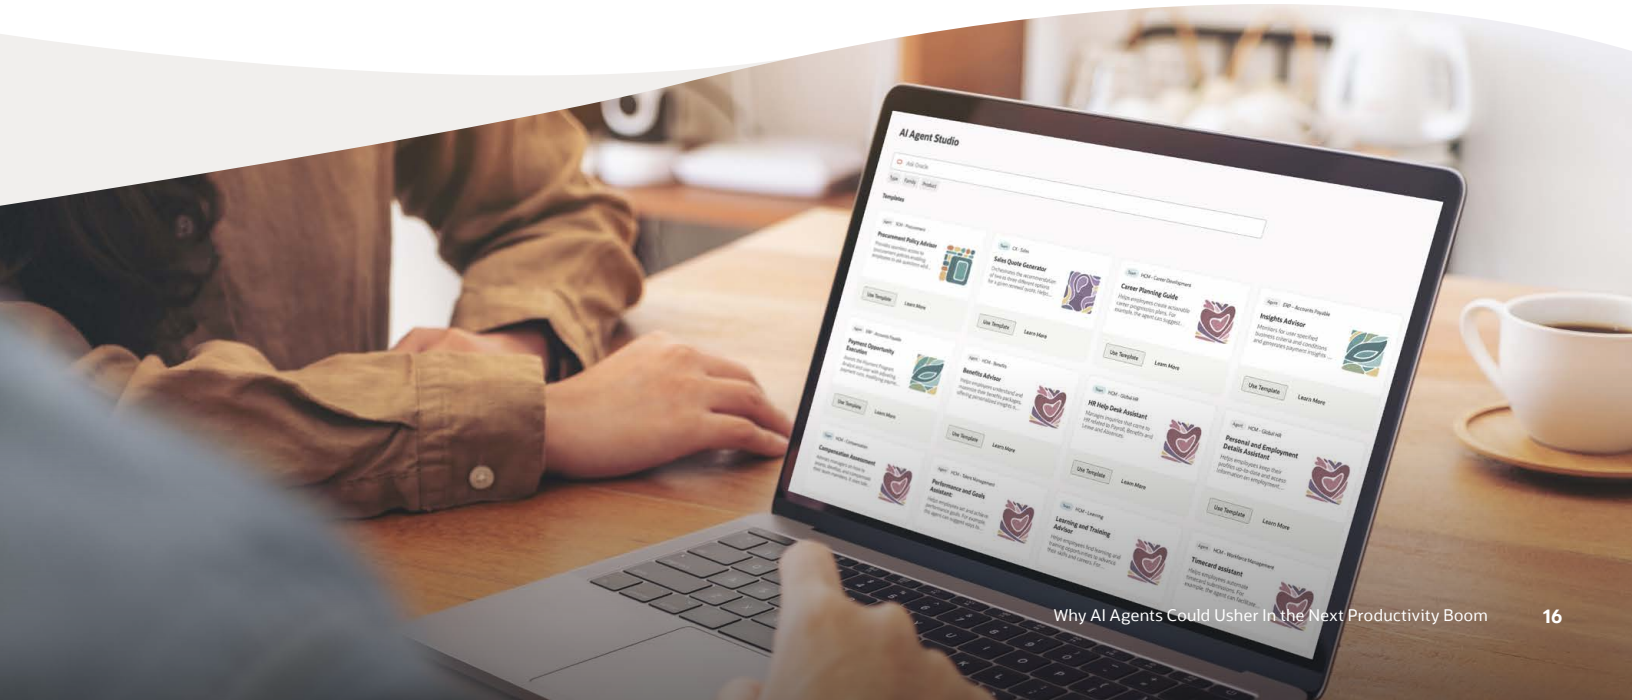

Agentic AI at work

Oracle's Fusion Applications, cloud infrastructure, and agent design studio are helping companies create AI agents that are designed to complete work faster in financial, supply chain, and other software.

Learn about deploying agents

Read about creating AI agents

Connect with us

Call +1.800.ORACLE1 or visit oracle.com

Outside North America, find your local office at oracle.com/contact

1 "From Potential to Profit: Closing the AI Impact Gap," Boston Consulting Group, Jan. 15, 2025.
2 "The next wave of AI: Demand and adoption," Barclays, Dec. 9, 2024.
3 "How to Marry Process Management and AI," Harvard Business Review, January-February 2025.
4 "Artificial Intelligence in the Knowledge Economy," Ide and Talamas, Feb. 25, 2025.
5 "Technology Vision 2025," Accenture, Jan. 7, 2025.
6 "The Agentic Evolution of Enterprise Applications," IDC, February 2025.
7 "User-Aligned Functions to Improve LLM-to-API Function-Calling Accuracy," Chan, Gentoro, Oct. 30, 2024.
8 "Measuring AI Uptake in the Workplace," Federal Reserve Board of Governors, Feb. 5, 2025.
9 "Usage of AI technologies increasing in EU enterprises," Eurostat, Jan. 23, 2025.
10 "Superagency in the workplace," McKinsey & Co., Jan. 28, 2025.
11 "Board Oversight of AI," Edelman Smithfield, Sept. 2024.
12 "The State of AI," McKinsey QuantumBlack, March 2025.
13 "Evident AI Leadership Report," Eviden Insights, Sept. 2024.