

ORACLE

The Rise of AI Agents: Unleashing Productivity and Innovation

Oracle Fusion AI





Table of Contents

Introduction	3
What are AI agents	4
AI agents for Oracle Fusion Applications	6
Examples of AI Agents for Oracle Fusion Applications	7
AI Agent use case examples	10
Summary	14

Introduction

Generative AI has sparked our imaginations and yielded invaluable benefits for early business adopters. Yet most organizations are only scratching the surface in learning what GenAI can do for them, and without the ability to securely use the company's own data and automate relevant tasks, applicability is limited for enterprises.

AI agents can be applied to address both problems. Working within established workflows, AI agents can leverage the power of large language models, interact with people, and consider internal data as they help solve complex problems alongside employees. Using AI agents within Oracle Fusion Applications can help transform the way work gets done by automating complex tasks and helping to provide valuable insights efficiently, extensibly, and consistently.

In this document, we'll discuss what agents are, provide examples of how they can help your organization, and explore the implications of their use within Oracle Fusion Applications.



What are AI agents?

AI agents combine large language models (LLMs) with other technologies and can be applied to accomplish complex tasks that previously could be done only by humans. Agents interact with their environments to gather data, determine the steps required to achieve a desired goal, and act on behalf of a role or persona. They can plan, use tools and data sources, make decisions with varying degrees of autonomy, and even work collaboratively with other AI agents.

LLMs are a core element of AI agent functionality and are what set them apart from the rules and machine-learning-based automation of the past.

Agents can be crafted to excel at achieving specific goals. That makes each one unique. However, there are traits they all share.

AI agents are:

Goal-oriented.

They assess and execute the tasks necessary to achieve a predetermined objective and can adapt to the environment in which they're working.

Autonomous.

AI agents can act on behalf of a user by, for example, invoking a tool, making a decision, initiating a process, or assisting another agent. While agents are designed to move work forward autonomously, humans are often in the loop to assess the agent's suggestions and guide, redirect, or overrule its recommended actions.

Specialized.

Agents adopt specific roles or personas and are designed to perform the tasks needed to achieve their prescribed goals.


Interactive.

Depending on its role, an agent may interact with humans in a conversational manner or communicate with other agents to request work, accept inputs, or send outputs to accomplish workflows.

Because a large language model is at the center of these interactions, AI agents can communicate with humans like a fellow human. They remember past interactions and can take feedback, iterate, and learn. LLM-powered AI agents can invoke logic to plan work and make high-quality, reasoned decisions.

AI agents are embedded in many applications. Examples range from simple automated scheduling and interactive Q&A apps to complex systems, like operating autonomous vehicles and advanced robotics. The design and capabilities of an AI agent depend heavily on the specific task it's meant to perform and the environment in which it will operate.

At the most complex end of the spectrum, agents can help tackle end-to-end strategic business processes by working with human employees who guide the AI agents toward their goal.



AI agents for Oracle Fusion Applications

For close to a decade, Oracle has developed and embedded AI functionality within Oracle Fusion Applications at no extra cost, across the entire suite, including CX, HCM, ERP, EPM, and SCM.

Oracle is not only leading with embedded enterprise AI functionality, it's forging ahead in deploying new, advanced AI technologies within Fusion Applications. AI agents are the latest example.


Oracle is expanding AI capabilities to advance beyond first generation GenAI and LLM systems, which relied on:

- Static LLM models trained on a specific moment-in-time data set and unaware of more recent events or information
- LLMs trained on publicly available data with no awareness of or access to specific business data
- LLMs with a “single request, single output” paradigm that lose the context of previous interactions

GenAI technology that can interact with its environment and users, remember how those interactions went, and be implemented to help call on other tools and agents to help when needed eliminates these challenges. The result is human-friendly technology with the power needed to accomplish complex tasks on behalf of and along with employees.

Examples of AI Agents for Oracle Fusion Applications

While Oracle has introduced the first set of RAG agents, the vision for future agents aligns around multiple agent types that will work together. Supervisory agents, conversational agents, functional agents, and utility agents cooperate to achieve desired outcomes. In a typical workflow, these agents interact, use tools, find necessary supporting data, make decisions, and unite to complete the task at hand.



Let's take a closer look at these classes of agents.

Conversational Agents

These are agents that interact with the outside world. In the case of enterprise applications, interactions are usually with humans, but they could be with another software program. In industrial settings, for example, conversational agents may interact with manufacturing equipment or Internet of Things devices.

Functional Agents

Functional agents, also called user-proxy agents, are most commonly associated with a particular organizational persona or role. Using a real-world example, you may encounter several “functional agents” when you go for your annual physical: The receptionist agent checks you in, and the nurse agent takes basic vitals such as your weight and blood pressure. Finally, you see the physician, the doctor who conducts a more detailed exam, assisted by an agent that summarizes the visit and generates necessary paperwork. Each of these agents performs specific subtasks, with specific expertise, using different tools, all communicating with one another as needed to accomplish a task.

Examples of functional agents include:

Hiring manager agent. Performs tasks including documenting requirements—for example, candidate skills and experience—which can be applied to assist with hiring decisions and reviewing job postings created by other GenAI systems for accuracy.

Field service agent. May help to provide information to technicians, to assist with automating tasks such as scheduling, diagnostics, and other decisions for more efficient field service workflows.

Receivables clerk agent. Assists with payment processing tasks, which may help with actions to improve cash flow, and assist with producing reports on receivables performance.

Customer support agent. Helps augment customer support functions with the potential to provide relevant information to human support agents or customers.

Supervisory Agents

Supervisors are the orchestra leaders among agents. These agents direct other agents and drive the planning and reasoning needed to achieve an objective. One type of supervisor is a user-proxy agent that makes decisions on whether to act on behalf of a human or connect with a person for human-in-the-loop feedback.



Utility agents

Agents that operate outside of common personas are referred to as utility agents, also called task-based agents. A utility agent is usually associated with a specific function and tool and is called on by other agents to perform a task, such as querying a database, sending an email, performing a calculation, or retrieving a document.

Utility agents deployed as part of a complex workflow usually act autonomously due to their low-risk functionality.

Examples include:

Copy generating agent. Helps to summarize a body of text or generates sample text to use as a starting point for longer communications.

Retrieval-augmented generation (RAG) agent. Assists with the retrieval of specific, up-to-date data necessary for an LLM to make a proper response to a prompt or carry out a task.

Skills enrichment agent. Uses HCM Dynamic Skills functionality to help suggest the skills needed to complete tasks, such as creating a job posting or assisting an employee with profile creation.

Database query agent. Helps to performs tasks related to data retrieval, such as making SQL queries.

Search agent. Helps to determine the optimal type of search, for example, a web or document search, and calls the appropriate tool to perform the task.

Coding agent. Writes code to perform a specific task using languages like HTML, Java, or Python.

Scheduler agent. Helps to schedule meetings with stakeholders to advance a project.

AI Agent use case example

HCM: Benefits Administrator

Joe, an employee working for a large financial services company, has an upcoming life event that has him wondering how he may be covered through his company-sponsored benefit plan. Through a conversational agent, a simple and familiar search field available through Oracle Cloud HCM, Joe can ask questions and receive personalized, accurate, and transparent answers.

For example, Joe can inquire through the agent what his insurance coverage is for an upcoming life event, such as the birth of his first child. Similar inquiries might be about emergency hospital coverage while on vacation, or comparisons of benefit plan coverage, such as deductibles, physician choices, and exclusions, considering Joe's expanding family.

Regardless of the question, the conversational agent passes the request to the supervisory agent, who creates the plan and determines the actions necessary to satisfy Joe's request. The supervisory agent may determine that an LLM should create the overall text composition of the response. However, specific knowledge regarding the company's benefit packages needs to be gathered. So, the supervisory agent will call on a RAG agent to fetch the appropriate benefits documentation. The supervisory agent may also direct an HR employee representative agent to retrieve employee information about Joe that may specify coverage possibilities and limits that pertain specifically to him. Finally, the supervisory agent will quality check the final response from the LLM for accuracy before forwarding it to the conversational agent. In the final response, specific sections of the policy document that pertain to Joe are highlighted.

Joe's experience can be enhanced further via the memory of agent interactions. Conversational agents can rely on their short-term memory to remember the conversation underway with Joe. For example, when provided a response about benefits coverage, Joe may ask, "I'm thinking about moving to Florida. Will this affect my coverage?" The agent, having retained the context of the conversation, will be able to help provide an appropriate and accurate response.



AI Agent use case example

SCM: Maintenance Agent

Imagine that Dania is a maintenance technician responsible for the maintenance of manufacturing equipment in her company's factory. An assembly line is down, and Dania is onsite to investigate the issue.

On her tablet, she dictates the symptoms that she sees with the equipment; that data is translated to text within the Oracle Fusion SCM application. She finds that a temperature gauge is indicating overheating, and the MG1000 press is inoperable. Along with describing the symptoms, Dania asks how she should proceed to fix the problem.

The supervisory agent receives the query from the conversational agent and creates a plan of action. Utility agents are called into action. The LLM determines potential causes, and a search agent is directed to retrieve MG1000 technical product documentation. An initial LLM response also includes additional steps that Dania should take to troubleshoot the problem, highlighting areas in the product documentation that show schematic diagrams for clarification.

She follows the troubleshooting steps outlined by the agent and concludes that a memory module is faulty and needs to be replaced. The response from the conversational agent asks for confirmation to proceed with a work order for a replacement part. She confirms, and the supervisory agent proceeds with the plan to replace the part.

A functional agent with the role of a procurement manager initiates the utility agents necessary to complete and authorize the order. One utility agent completes the purchase order, and another sends the electronic order to the vendor. Another utility agent sends a detailed email to Dania confirming the order with high-priority shipping.

In this example, various classes of agents coordinate actions, make decisions—with human confirmation when necessary—and fulfill multistep workflows.



AI Agent use case example


ERP: Payables agent

Within the procure-to-pay cycle, the accounts payable process involves the steps necessary to process and pay invoices from vendors and suppliers. Agents can help to improve the efficiency of this process within Oracle Cloud ERP.

An accounts payable workflow may be triggered without the use of a conversational agent; it can be initiated through a user-proxy agent or by a functional agent assuming the role of a payables clerk. Each day, the clerk agent triggers the utility agents required to gather invoices from various sources and prepare them for processing. Supervisory agents help to plan the automated workflow, calling utility agents that can predict and populate code combinations for non-purchase-order (PO) invoices. Or, for invoices associated with POs, a utility agent may direct the steps necessary for invoice-line to PO-line matching. Other utility agents, acting under the direction of functional agents, automate the invoice approval process and then, in turn, initiate the steps necessary for payment.

For exceptions that require human intervention, a utility agent can route the invoice and notify the appropriate manager that further validation is required.

Agents automate actions, make low-risk decisions, and involve humans only when necessary. Now, a multistep process that previously took days to complete may happen in hours, with no intervention.



AI Agent use case example

CX: Automated Customer Service Agent

Today, Alex, a service agent, and her colleague Adam, a field service technician, work together as part of a larger service team to resolve customer issues. They must navigate complexities like changing customer expectations, Service Level Agreements (SLAs), and supply chain shortages, while managing high volumes of cases and maintaining operational efficiency.

Instead, AI Agents in Oracle Service can be used to help automate large portions of the service workflow. For example, when a piece of IoT-connected equipment reports a malfunction, an AI Agent is designed to immediately analyze the issue, determine the likely root cause (e.g., a faulty sensor), generate an action plan, and then begin execution of the action plan. This plan may include ordering replacement parts, creating work orders, or scheduling field technicians.

When anomalies occur, the AI Agent can be used to flag the issue and propose a solution to Alex, within her workspace. For instance, if a required part is backordered, the AI Agent can be used to recommend sourcing the part from a third-party supplier but will seek Alex's approval before proceeding.

In this example, the automated service workflow could look like this:

1. Automated Service AI Agent diagnoses the equipment issue and prepares an action plan.
2. AI automatically creates a service request, orders parts, and schedules the repair.
3. If a process exception is detected (e.g., part shortage), the AI Agent alerts Alex (service agent), recommends a potential solution, and requests approval to move forward.
4. Once approved, the AI Agent continues executing the plan, which may include updating the customer, managing supplier orders, and scheduling Adam's visit to the customer site.
5. After Adam completes the repair, any notes are logged to help inform future incidents.

By automating routine tasks and handling more complex decisions through collaboration with Alex and Adam, Oracle Service helps service teams reduce resolution times, meet SLA targets, and increase customer satisfaction.



Summary

Accelerated advancements in the capabilities of AI agents have the potential to fundamentally change how humans—and other machines—work with enterprise applications across HCM, ERP, CX, and SCM. Agents can be implemented to help improve on the relatively recent introduction of generative AI, and advancements in agent technology will no doubt accelerate too.

As of this writing, agent technology is in its introductory phases. The potential for agents to act autonomously to manage entire workflows is enticing. However, responsible use of AI still requires controls, safeguards, and human oversight.

Oracle is leading the way in developing and integrating AI agents within Fusion Applications in a manner that provides tight control over data security and governance. We're helping our customers fundamentally transform the way work gets done by introducing more:

Efficiency

Automate repetitive tasks and allow your team to focus on strategic activities.

Scalability

Adapt to growing and changing business needs without increasing staffing.

Consistency

Enable reliable and uniform quality and performance across all interactions and tasks.

Explore new ways of working

Learn how Oracle AI and Fusion Applications can help improve productivity and efficiency by working alongside your employees to help accomplish complex tasks and automate workflows.

[Learn more](#)

Connect with us

Call +972-3-9273700 or visit oracle.com/il-en/

Outside the Israel, find your local office at oracle.com/emea/corporate/contact/

Copyright © 2025, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

