

## EXPLAINER · ADVANCED AUTOMATIONS AND AI AGENTS

# An AI agent is software that takes actions toward a goal, and it needs a named owner and firm limits.

Published 18 October 2026    3 min read




Photo: Homa Appliances, Unsplash

## IN BRIEF

An AI agent acts toward a goal by planning steps and using other tools and systems, which a chatbot that waits for each prompt does not do.

---

Government guidance in Canada and the UK points organizations to narrow, low-risk tasks and tight limits on what an agent can reach.

---

Before an agent touches real systems, name who owns it, who approves its access, who watches it and who can switch it off.

## What an AI agent is

The Government of Canada's guide on agentic AI describes an AI agent as a system that can perceive and act on its environment, often autonomously, to achieve specific goals. The same guide says agentic AI is defined more by what the system does than by what it produces, which is a useful test for a leader.

A tool that drafts a reply produces content. A tool that reads the request, finds the record, updates it and sends the reply is acting.

## How it differs from a chatbot and from automation

OECD.AI says that earlier generative AI tools mainly respond when prompted, while an agentic system can plan the steps and draw on other tools and systems to reach a goal. A person sets the goal and the boundaries instead of directing each step.

Rule-based automation is a different thing again. The European Commission's 2025 guidelines on the AI Act describe basic data processing as software that follows predefined, explicit instructions and works from fixed human-programmed rules, and place it outside the Act's definition of an AI system.

The guidelines do not discuss agents. In our reading, a rule-based automation repeats the same steps every time, and an agent chooses its steps within the limits it is given.

## What organizations use agents for

Public bodies advise caution about where to start. The Government of Canada's guide says agentic AI is best used in tightly scoped internal workflows with clear boundaries and limited permissions. It sets four levels of autonomy, using meeting scheduling as its example: software that suggests times, a system that drafts an invitation for approval, one that books the meeting on delegated permissions and one that reschedules within set limits and escalates exceptions.

The UK National Cyber Security Centre, summarizing joint guidance written with international partners, says agentic AI is likely to offer significant benefits in many scenarios, particularly where tasks are repetitive, well understood and low risk. It advises starting with tightly bounded pilots.

Practice appears to follow that advice: in OECD interviews with 25 organizations, none reported deploying agentic systems with unrestricted autonomy. Human approval for high-stakes or irreversible actions, such as payments and deleting data, was typical, although the threshold varied. The OECD calls the findings illustrative and not representative.

## What oversight an agent needs

The same OECD paper reports that interviewees said agents may take actions that are not fully predictable and may vary from one run to another, and that errors may surface only over time.

The UK National Cyber Security Centre puts the readiness test plainly: if you cannot understand, monitor or contain an agent's actions, it is not ready for deployment. It adds that people stay accountable for the decision to deploy, the access granted and the consequences. Who owns the agent, approves its access, monitors it, reviews incidents and can stop it should be settled before it connects to real systems or data.

The Government of Canada's guide, written for federal departments and agencies, adds that accountability rests with a designated human owner. It also asks for human checkpoints where an action changes a system. It asks for a log of the agent's actions that the agent cannot alter, and for a way to pause or disable the agent from outside it.

## What this means for you

Start with one task that is repetitive, well understood and low risk. Decide what the agent may read, change and send before you build anything. Give it the minimum access for the shortest time, and require a person's approval for payments, deletions and anything hard to reverse.

Name an owner, an approver of access, a monitor and someone with authority to switch the agent off. Write the names down before the agent connects to live systems, keep a log of what it does and review it on a set schedule.

Our advanced automations and AI agents service helps leaders choose a first task and set these limits. If you would like to start, request a briefing.

---

## SOURCES

Government of Canada [Guide on the Use of Agentic Artificial Intelligence](#) 2026

OECD.AI [Governing with agentic AI: when machines act on government's behalf](#) 2026

European Commission [Commission Guidelines on the definition of an artificial intelligence system established by Regulation \(EU\) 2024/1689 \(AI Act\), C\(2025\) 5053 final](#) 2025

UK National Cyber Security Centre [Thinking carefully before adopting agentic AI](#) 2026

OECD [Agentic AI in organisations: Early insights from practitioner interviews, OECD Artificial Intelligence Papers No. 65](#) 2026

---

## DISCUSS THIS WITH US

To discuss this with Consier Technologies, visit [consiertech.com/contact/](https://consiertech.com/contact/).