

AI agents that act on your behalf: Set the right limits

Adaptive module

~20 min

AI agents no longer just recommend: they make payments, reply to customers, change systems, invite candidates to interviews, and commit the company before any human review. Delegation training was not designed for this kind of delegate. This module applies the logic of delegation of authority to agents: a six-clause mandate, autonomy proportionate to the risk, forbidden actions made impossible through configuration, and accountability that stays human, from the organization to the team member.

— LEARNING OBJECTIVES

- 01 Distinguish delegating a task from delegating the power to act, and name the four powers given to agents: money, communication, systems, and people.
- 02 Build a six-clause mandate: outcome, playing field, allowed and forbidden actions, thresholds, stop points, duration and record.
- 03 Match the agent's autonomy to the cost and reversibility of an error, and make forbidden actions impossible through configuration.
- 04 Revise or withdraw a mandate using the activity log, and explain the accountability of the organization, manager, and team member.

— WHO IT'S FOR

Team members who want to give an AI agent real power, and managers who want to set the framework for it.

- **Prerequisite:** none

— HOW THE MODULE WORKS

Self-paced, with practice tailored to each learner.

- 1 **Self-assessment:** on 9 dimensions, to gauge current habits and guide the practice.
- 2 **Expert video:** 10 minutes, with 3 pauses for reflection.
- 3 **Practice scenarios:** the first 2 chosen from the self-assessment, then 16 optional ones. Every answer gets feedback: what it leads to, and one way to improve.
- 4 **Dated commitment:** one action to try, written by the learner or picked from the suggestions.
- 5 **Memo to download:** key takeaways, questions, commitment, personal notes and checklists.

— A SAMPLE PRACTICE SCENARIO

For three months, Élise has approved every batch of corrections her AI agent makes to product records: labels, units, customs codes. Élise has never changed anything in the agent's proposals. Her manager asks her how far to let the agent correct on its own. How can Élise decide?

— ASSESSMENT AND COMPLETION

Pass mark: 80%

The score takes into account your answers to the practice scenarios, your commitment and, where relevant, any additional learning activities. Passing requires a commitment and a score of at least 80%.

METHODS & MODELS

- The four powers that can be delegated to an agent: committing money, speaking for the organization, changing systems, taking action affecting people – Afthonios method
- The mandate in six clauses: result, playing field, allowed and forbidden actions, thresholds, stop points, duration and audit trail – Afthonios method
- Autonomy proportionate to three questions: cost of an error, reversibility, consequences outside the organization
- Reviewing the mandate based on the activity log: widen, tighten, or withdraw

SOURCES & REFERENCES

- Gartner (March 2025): by 2029, 80% of routine customer service requests will be resolved by agentic AI without human involvement
- OpenClaw incident (Meta, February 2026): an agent deletes more than 200 emails after losing the confirmation instruction
- Moffatt v. Air Canada (2024): a company is accountable for what its chatbot says
- AI Act (EU Regulation 2024/1689), Annex III: recruitment is a high-risk use, subject to human oversight