

COURSE 04 · BUSINESS TRACK

AI Automation

From process maps to agent-built automations a business can run

3 DAYS

21 CONTACT HOURS

HANDS-ON

A practical course for people who automate business processes: choosing what to automate on evidence, building integrations with coding agents rather than by hand, and running the result with monitoring, cost control and a handover a team can own.

AUDIENCE

Operations leads, automation specialists and consultants, analysts, technical product managers and engineers moving into automation.

DELIVERY

One real or realistic business process is mapped, scored, automated, hardened and handed over across the three days.

CAPSTONE

One automated business process with a measured baseline, run cost, error handling and a handover pack.

Feasibility verdict: Feasible when every team automates one process end to end against sandbox systems. Integrating with live company systems, or comparing every automation platform hands-on, would not fit into three days.

Version 1.0 · October 2026 · dinu.at/teaching/ai-automation

Course overview

How to find the work in a business that is worth automating, have AI agents build the automation, and show afterwards that it saves more than it costs. One real or realistic business process is mapped, scored, automated, hardened and handed over across the three days.

Learning outcomes

- Map a business process into tasks and score each task by frequency, time, cost of errors, data access and how easily its output can be checked.
- Write a business case for an automation: hours saved, the value of that time, build and run cost, maintenance and break-even, with the assumptions stated.
- Decide when a drag-and-drop builder (Zapier, Make, n8n, Power Automate), a coding agent (Claude Code, Codex), a small internal app or no automation at all is the right choice.
- Direct a coding agent to build integrations against APIs, webhooks and MCP connectors from a written specification, and review what it produced.
- Design approval points, idempotent steps, retries and fallbacks, so that a failed run is noticed and can be repeated safely.
- Produce automation results beyond text, such as filled documents, slide decks, images and short videos, and check them before they reach a customer.
- Monitor an automation in production through success rate, exceptions, cost per run and drift, with a named owner and an escalation path.
- Scope, price and hand over automation work with a proposal that states measurable outcomes, data-protection duties under the GDPR and the EU AI Act obligations that apply.

Realistic depth within three days

IMPLEMENT	PRACTISE	SURVEY / DEMO
• Process audit and automation scorecard • End-to-end automation built with a coding agent • Monitoring, alerts and a cost report	• Agent-written API and MCP integrations • Approval checkpoints, retries and idempotency • Proposal, pricing and handover pack	• Enterprise RPA and process-mining suites • Every drag-and-drop platform in depth • Custom model training or fine-tuning

Conditions that make the schedule feasible

- Each team brings one process from its own work or picks one from a prepared catalogue (lead intake, invoice handling, support triage, recurring reports).
- All labs run against sandbox APIs, a mock CRM, a ticketing system and synthetic data; no lab touches production credentials or live customer data.
- One coding agent (Claude Code or Codex) and one visual builder (n8n) are used hands-on; other platforms are discussed, not installed.
- A prepared starter repository provides connectors, test data and a deployment target, so lab time goes to the automation rather than to setup.
- Roughly 10 hours of guided building, 8 hours of concepts and demonstrations, and 3 hours of business cases, review and presentations.

Recommended participant profile

Primary audience. Process and operations owners, automation specialists, freelance and agency automation consultants, business analysts, technical product managers and engineers who are asked to automate business work.

Prerequisites. Comfort with spreadsheets, web tools and reading short scripts. No programming background is required, but participants should be willing to read the code an agent writes and to work in a terminal with guidance. A laptop on which the starter environment can be installed.

Class size. 10–18 participants per instructor; a teaching assistant above 14, or when most participants have never used a terminal.

Day 1: Choosing work worth automating

Decide what to automate on evidence before building anything.

Time	Module	Content and activity
09:00–09:30	Orientation	Automation as the use of technology to reduce labour; what changed once agents started writing integrations; the course process and the capstone.
09:30–10:30	What agents automate today	Live demonstration: a coding agent builds a working integration from a one-paragraph request. Where it succeeds, where it fails, and why checking the result is now most of the work.
10:30–10:45	Break	Coffee and questions.
10:45–12:15	Process discovery LAB	Interview a process owner, map the process, split it into tasks and record volumes, durations, error rates and the systems each task touches.
12:15–13:15	Lunch	Lunch break.
13:15–14:00	Scoring tasks	Frequency, time per case, cost of an error, data access, how checkable the output is and how much judgement it needs. Why digitised, frequent and verifiable tasks go first.
14:00–15:00	The economics of an automation	Hours saved and what they are worth, build cost, cost per run, maintenance, the cost of failures and break-even. Why most of the value usually sits in a few tasks.
15:00–15:15	Break	Coffee and questions.
15:15–16:15	Choosing the tool	Visual builders, coding agents, small internal web apps, off-the-shelf software, or leaving the task with people. Lock-in, ownership and who maintains the result.
16:15–17:30	Lab 1: business case LAB	Score the task inventory, choose the automation target and write a one-page business case with its assumptions stated.
DAY OUTPUT A process map, a scored task inventory and a one-page business case for the process the team will automate.		

Day 2: Building automations with agents

Turn the chosen process into a working automation that an agent builds and a person reviews.

Time	Module	Content and activity
09:00–09:30	Writing the specification	Inputs, outputs, systems, edge cases, approval points and acceptance tests, written so that an agent can build from them and a reviewer can check against them.
09:30–10:30	Integration building blocks	APIs, authentication, webhooks, schedules, queues, MCP connectors and skills. Structured outputs and schema validation between steps.
10:30–10:45	Break	Coffee and questions.
10:45–12:15	Lab 2: agent-built integration LAB	Direct a coding agent to build the first half of the process against the sandbox (for example intake, enrichment and classification), then review and test its code.
12:15–13:15	Lunch	Lunch break.
13:15–14:00	People in the loop	Approval checkpoints, confidence thresholds, escalation and audit trails, and how to keep an approver from becoming a rubber stamp.
14:00–15:00	Reliability	Idempotent steps, retries, timeouts, dead-letter queues, partial failures and re-running a single case safely.
15:00–15:15	Break	Coffee and questions.
15:15–16:15	Outputs beyond text	Generating documents, slide decks, images and short videos as automation results; templates and brand rules; checking generated media before it is sent.
16:15–17:30	Lab 3: completing the flow LAB	Add the second half (for example a drafted proposal and a CRM or ticket update), the approval step and tests for every edge case in the specification.

DAY OUTPUT

A working end-to-end automation in the sandbox, with its specification, tests and an approval step.

Day 3: Running automations in a business

Turn a working automation into something a team can own, measure and pay for.

Time	Module	Content and activity
09:00–09:30	Monitoring and evaluation	Success rates, exception queues, sampling outputs for review, cost per run, and drift when inputs or upstream systems change.
09:30–10:30	Data protection and regulation	GDPR roles and data minimisation, processing agreements, personal data in logs, and the EU AI Act obligations that apply to organisations deploying AI systems.
10:30–10:45	Break	Coffee and questions.
10:45–12:15	Lab 4: monitoring and cost LAB	Add logging, alerts and a cost report to the capstone and run it on a batch of test cases, including deliberately broken ones.
12:15–13:15	Lunch	Lunch break.
13:15–14:00	Scoping and pricing	Fixed-scope, time-based and value-based pricing; a proposal that names measurable outcomes, exclusions and maintenance; why a small scope that ships beats a large one that doesn't.
14:00–15:00	Handover and adoption	Runbooks, ownership, access, change requests, training the people whose work changes, and measuring the result after launch.
15:00–15:15	Break	Coffee and questions.
15:15–16:15	Capstone completion CAPSTONE	Finish the handover pack and the measured before-and-after comparison.
16:15–17:30	Demonstrations and review CAPSTONE	Team demonstrations against the business case, a cost and failure review, and next-step recommendations.
DAY OUTPUT A monitored automation with a cost report, a handover pack and a proposal or internal decision paper.		

Labs and capstone

Lab 1: Process audit	Process map, scored task inventory and a one-page business case.
Lab 2: Agent-built integration	Specification, agent-written integration code, review notes and tests.
Lab 3: Approval and reliability	The complete flow with an approval step, retries, idempotent steps and edge-case tests.
Lab 4: Operations	Monitoring, alerts, a cost report, the handover pack and a proposal.

Capstone specification

- A process map and scored task inventory that explain why this process was chosen.
- A baseline measured or estimated before automation: volume, time per case, error rate and cost.
- An automation built mainly by a coding agent from a written specification, with the specification and the review notes included.
- At least one human approval point, and documented behaviour for every failure the specification names.
- Idempotent steps, so any failed case can be re-run without duplicate side effects.
- A cost-per-run figure and a break-even estimate based on the baseline.
- A handover pack: runbook, owner, access list, monitoring view and data-protection notes.

Tools and materials

Process mapping	Mapping template and scoring sheet Most automation mistakes happen before anything is built; a simple scorecard makes the choice explicit and reviewable.
Coding agent	Claude Code or Codex Writes integrations from a specification, which is how most new automations get built now. One is used hands-on, the other is compared.
Visual builder	n8n (self-hosted) Shows where a visual workflow is still the clearer choice, and gives a fair comparison with agent-written code.
Connectors	REST APIs, webhooks and MCP servers The common ground every platform builds on; MCP lets agents and tools share the same connectors.
Sandbox systems	Mock CRM, ticketing system and mailbox Realistic side effects without touching production data.
Generated outputs	Document and slide templates, image and video generation APIs More and more automation results are documents and media, not just fields in a database.
Operations	Structured logs, alerts and a cost report An automation nobody watches fails silently, and cost per run decides whether it pays.

Discussed, not taught hands-on

- Zapier, Make and Power Automate are discussed and compared but not installed; n8n stands in for the visual builders.

- Enterprise RPA and process-mining suites are surveyed for context. They suit large, stable process estates more than the fast-changing processes this course targets.
- Lead generation and sales for automation agencies are covered only as far as scoping and pricing need them.

The course teaches working proficiency in choosing, building and running automations. It does not certify compliance, and a capstone built against sandbox systems still needs a security review and a data-protection assessment before it touches production.