


FOR PRODUCT MANAGERS · FOUR WEEKS

The Agentsable syllabus.

Four weeks of live agent runs, plain text builds, and hands on browser labs.



Run live agents.

Master Raw Agent Control (RAC).

Become Agentsable.

Go from running your first live agent in minute 15 to architecting a production ready multi agent workflow, with zero coding and zero framework lock in.

One rhythm. Every week.

Designed for busy PMs. Missed the recordings? The live concept review and RAC build bring you up to speed before you take control.

01

ON DEMAND LECTURES

Recorded videos unlocked weekly.

Short, focused video breakdowns released after each weekly session so you can explore the core concepts at your own pace.

02

WEEKLY LIVE Q&A

Fast concept recap & live Q&A on Zoom.

We kick off each live session with a clear review of the recorded material and answer your specific questions.

03

LIVE RAC BUILD SESSION

Watch a live agent system built in plain text.

We build a complete, working example live on screen using the RAC system. Even if you missed the recorded videos, the Q&A and Live RAC Build give you 100% of what you need.

04

HANDS ON BROWSER LAB

Take control by editing plain text.

Immediately after the live build, you get a ready to run version of that exact system in your browser lab and implement it yourself purely by editing the plain text files and testing execution, with zero local installations or IT permissions required.

WEEK 01 · AGENT BRAINS

From prompt to agent

Defining a structured plain text agent

Stop writing fragile chat prompts. Start architecting structured agent brains.

WHAT YOU'LL EXPLORE

01 First 15 minutes: run a live agent

Trigger a live agent in the browser lab and see firsthand why a standard chat prompt breaks when real world inputs get messy.

02 The anatomy of a RAC agent

Move from unstructured prompting to a structured plain text RAC definition. Map the agent's Role, Business Context, Decision Logic, and Output Contract.

LIVE LAB CHALLENGE

Take over the agent from the live build that is misclassifying edge cases. Fix its reasoning behavior by refining its plain text RAC rules.

WEEK 02 · TOOLS & MEMORY

Tools, MCP & memory

Connecting external tools, Model Context Protocol (MCP), memory, and multi agent handoffs

Orchestrate multi agent teamwork in plain text.

WHAT YOU'LL EXPLORE

01 Plain text tool & MCP definition

Demystify external tool calling and the Model Context Protocol (MCP). Define which tools and data sources an agent can access, and the exact rules for calling them, in readable text.

02 Memory & context window control

Define what an agent remembers across steps, what gets pruned, and how to prevent context bloat before it confuses the model.

03 Multi agent handoffs

Wire specialist agents together by defining strict context handoffs in plain text.

LIVE LAB CHALLENGE

Take the multi agent workflow from the live build and wire its tools, memory rules, and agent to agent handoffs in plain text so the team completes a multi step task end to end.

WEEK 03 · FAILURE PATTERNS

When agents fail

Failure patterns of agentic workflows

Design reliable AI systems by seeing where live agents break.

WHAT YOU'LL EXPLORE

01 Opening the black box

Read live execution traces in real time to diagnose why an agent made a bad decision instead of guessing.

02 Infinite tool calling loops

Intentionally trigger and inspect an agent receiving ambiguous data and repeating the same tool call in circles.

03 Context decay & instruction dropout

See how long workflows can cause an agent to forget its original rules mid task.

04 Silent hallucinations & broken formats

Inspect what happens when an agent invents parameters or breaks the format expected by the next step.

05 Cascading multi agent errors

Trace how one agent's small mistake snowballs across an entire team of agents.

LIVE LAB CHALLENGE

Inspect a broken multi agent workflow trace from the live build. Diagnose the exact failure pattern and stop the loop by adjusting the plain text definitions.

WEEK 04 · PRODUCTION CONTROL

Production control

Guardrails, evaluators & human oversight

Turn a fragile workflow into a controlled system with explicit boundaries and quality checks.

WHAT YOU'LL EXPLORE

01 Hard plain text guardrails

Move beyond polite prompt suggestions. Set maximum reasoning steps, blocked actions, and strict execution rules outside the prompt.

02 Pass/repeat quality gates: LLM as a Judge

Build automated evaluator gates that inspect an agent's output against your quality rubric and send bad work back for a retry before it moves forward.

03 Human in the Loop systems

Design uncertainty triggers and approval gates so a workflow pauses and routes high stakes or ambiguous decisions directly to a human PM.

04 Production codebase handoff

Package your tested plain text RAC definitions so your engineers or AI coding tools can implement the exact workflow directly into your company's existing codebase, with zero framework lock in.

LIVE LAB CHALLENGE

Add plain text guardrails, an LLM as a Judge pass or repeat gate, and a Human in the Loop escalation rule to the live build example. Practise the controls needed for production readiness.