

THE AI MARKETING-OPERATOR COHORT · BY TALENT SCOUT

THE FULL CURRICULUM.

The four-week arc, broken down to the session level.
Every module, every deliverable, every tool — mapped out.

FORMAT

8 live sessions over 4 weeks · Tuesday & Thursday, 2 hours each
Monday pre-read video (15–20 min) · Independent build time between sessions

PLATFORM

Built entirely on Claude Cowork — one foundation, every agent, every automation, every artifact

INSTRUCTOR

Felipe Oviedo — AI Solutions Director, Talent Scout · 200+ AI systems deployed

4

WEEKS

8

SESSIONS

16

TRAINING HOURS

85%

HANDS-ON

gettalentscout.com/ai-cohort · \$1,900 per seat · Cohorts run monthly

FOUNDATIONS + SETUP

AI Landscape + 3 Workspaces + MCP Intro

The AI landscape mapped. Three specialized workspaces built from scratch. Claude Cowork Projects deep training. MCP connectors overview.

SESSION 1 — AI LANDSCAPE & CLAUDE COWORK SETUP

- Claude vs ChatGPT vs Gemini vs Perplexity, honest comparison module
- Three specialized workspaces: Research, Content, Operations
- Claude Cowork Projects deep training (architecture, naming, routing)
- MCP connector overview: what connects, what doesn't, what's coming

SESSION 2 — KNOWLEDGE LAYER & PROJECT ARCHITECTURE

- Obsidian vault setup as a personal knowledge layer
- Template: client workspace architecture (naming, structure, routing)
- First project architecture documented end to end
- Knowledge compounding system that works across clients and projects

SHIPS THIS WEEK

3 connected workspaces (Research, Content, Operations) + documented project architecture + Obsidian vault initialized

FIRST MARKETING AGENT

RCTFX Framework + System Instructions + Agent Deployment

The RCTFX framework for production-grade prompting. System instruction templates built by hand. First agent deployed on a real recurring task.

SESSION 3 — PRODUCTION-GRADE PROMPTING

- The RCTFX framework: Role, Context, Task, Format, eXamples
- 5 system instruction templates built by hand, tested on real data
- Natural-conversation playbook for everyday AI work
- Prompt debugging: when Claude drifts, how to fix it

SESSION 4 — FIRST AGENT DEPLOYMENT

- Pick one real recurring task from their current role
- Build the agent end to end: system instructions, context, triggers
- Deploy it on real client data for a task they actually do every week
- Live proof their AI is working for them, not the other way around

SHIPS THIS WEEK

1 live agent deployed on a real recurring task + 5 production-grade system instruction templates + natural-conversation playbook

AUTOMATION + KNOWLEDGE

4 Automation Types + Knowledge Bases + MCP Deep Dive

The four automation types every operator needs. Two automations deployed. Two knowledge bases loaded and verified on-brand. MCP deep dive.

SESSION 5 — AUTOMATION ARCHITECTURE

- The four automation types: scheduled tasks, custom skills, live artifacts, multi-step workflows
- When to use which type (decision framework)
- Build and deploy automation #1: a scheduled task on real data
- Build and deploy automation #2: a custom skill or live artifact

SESSION 6 — KNOWLEDGE BASES & MCP ECOSYSTEM

- Two knowledge bases loaded: brand voice + client context + historical performance
- Verification loop: test that the AI stays on-brand with real prompts
- MCP ecosystem deep dive: marketing-grade connectors
- Connecting external tools to Claude Cowork via MCP

SHIPS THIS WEEK

2 automations running on schedule + 2 knowledge bases loaded and verified on-brand + MCP stack live

FULL STACK + SHOW & TELL

Marketing Operator Stack + Live Artifact + Public Demos

The six-layer Marketing Operator Stack assembled. A live HTML artifact deployed to a public URL. Public Show & Tell: three-minute live demos.

SESSION 7 — FULL STACK ASSEMBLY

- Full 6-layer Marketing Operator Stack assembled with final checklist
- Build front-end HTML artifact live with Claude Cowork
- Deploy artifact to live URL (GitHub Pages or equivalent)
- 3 governance rules written and implemented

SESSION 8 — SHOW & TELL

- Three-minute live demos in front of an invited audience (including leadership)
- Operational playbook finalized: what runs, what's manual, what's next
- Peer feedback and cross-team insights
- Post-cohort roadmap: what to build in month 2 and beyond

SHIPS THIS WEEK

Complete 6-layer Marketing Operator Stack + live HTML artifact at a public URL + governance playbook + operational playbook

GETTING STARTED

BEFORE THE FIRST SESSION

Nothing exotic. If your team member can already use Claude, Google Workspace, and a Chrome browser, they're 90% of the way there.

TOOLS TO HAVE READY

- Active Claude account (Pro or Team tier for Cowork access)
- Google Workspace or equivalent (Gmail, Drive, Docs)
- Obsidian (free download)
- Two-monitor setup recommended (one for class, one for build)

TIME COMMITMENT

- 15–20 min pre-session video each Monday
- 4 hrs/week live sessions (Tue & Thu, 2 hrs each)
- 1–2 hrs independent build time between sessions

TOTAL: 6–8 HOURS/WEEK × 4 WEEKS

WHAT TO BRING TO CLASS

- A real recurring task in their current role they want to automate
- Real client data: brand voice, briefs, past performance, campaign context
- Willingness to demo their build at the Week 4 Show & Tell

THE WEEKLY RHYTHM

MONDAY

Pre-read & prep — trainer video (15–20 min)

TUESDAY

2-hour live session + interactive workshop

WEDNESDAY

Apply & build independently on real work

THURSDAY

2-hour live session + interactive workshop

ALL 8 DELIVERABLES

3–5 specialized Claude Cowork workspaces

5–10 production-grade system instruction templates

1 deployed marketing agent on a real recurring task

2 active automations running on schedule

2 knowledge bases verified on-brand

Personal Obsidian vault for compounding knowledge

Live HTML artifact deployed to a public URL

Complete governance playbook

READY TO BUILD AN **AI-FLUENT TEAM?**

20 seats per cohort · \$1,900 per seat · Group discounts for 5+

RESERVE A SEAT →

gettalentscout.com/ai-cohort