

A PRACTICAL GUIDE

The AI Orchestrator Playbook

A framework for directing several specialized agents instead of running one tool for everything.

INTRODUCTION

You're Already Doing More Than One Job

What used to be "using an AI tool" is quietly turning into "coordinating several". Even a single freelancer today might have one agent drafting, another researching, another handling admin, without ever formally deciding they built a team, it happens by accretion, one convenient tab at a time, not by an intentional choice.

This guide is about naming that shift and giving it structure: becoming an orchestrator on purpose instead of by accident. It isn't about which specific agents to use, that's a decision for another day, it's about the decision layer that sits above them: who handles what, and who decides.

SECTION 1

The Cost of Running One Tool for Everything

GENERALIST DRAG

A single, all-purpose tool averages out across tasks: good enough at drafting, adequate at research, mediocre at chasing down details, because none of those is what it was actually optimized for.

INVISIBLE SWITCHING COST

Even people who do use different tools for different jobs often decide which tool handles what in the moment, re-deciding the same routing question over and over. That's its own quiet tax.

The fix isn't collecting more tools. It's separating two jobs that got merged into one habit: doing the work, and deciding who, or what, does the work.

SECTION 2

The Orchestrator Layer

THE HANDOFF DECISION

OBJECTIVE Decide, before starting, which agent (or which mode of one agent) is the right fit for this specific task, not the one that happens to already be open.

METRIC A one-line note of which agent handled which part of a finished piece of work, so the routing becomes visible instead of automatic.

BOUNDARY Never let "it's already open" be the reason a task goes to a given agent.

A SIMPLE ROLE MAP

Input agents. The ones that take in and organize: research, transcription, first-pass drafts. Good at breadth, not at final judgment.

Output agents. The ones that produce a finished thing: writing, formatting, scheduling. Good at execution once direction is clear.

Decision points, not agents. The moments where nothing gets handed off, because the choice, what to say yes to, what to charge, what to publish, is still yours by design.

PROMPT · MAP YOUR OWN STACK

Here is a list of the AI tools and agents I currently use in my work: [list them]. For each one, tell me whether it's mainly an input tool (gathering and organizing), an output tool (producing a finished deliverable), or overlapping with another tool on the list. Flag any task I'm currently handing to whichever tool is already open, rather than to the best fit for that task.

The point isn't to have more agents. It's knowing, without re-deciding it every time, which one earns a given task.

Directing several specialized agents well beats running one generalist tool for everything.

The orchestrator role is a decision, not an accident.

Once it's explicit, everything downstream gets faster to route.

**Follow Martina La Porta for more on navigating an AI-era career,
or subscribe to AI's Other Shore on Substack.**