

The Orchestrator's Field Workbook

Every exercise from Orchestrating Intelligence, laid out to be filled in.

Anand Padia · orchestratingintelligence — free companion edition

A book you only read changes what you know. A book you *work* changes what you do. This workbook turns the ideas into a set of tools you can run on your own organization — alone at your desk or with your team in a room. Photocopy the pages, or rebuild them on a whiteboard. The point is to fill them in.

Use them in roughly this order: tag your processes, map one loop, score yourself on the Audit, check a project against the Stack, build a scorecard, then run a thirty-day sprint. The consolidated action list at the end gathers all thirty-five "Put It to Work" moves in one place.

Tool 1 — The Watch-or-Murmuration Tagger

List your five most critical processes. For each, decide whether it needs a *guarantee* (a watch) or needs to *adapt* (a murmuration), using two questions: how much does the world around it shift, and how bad is a single mistake? Defend any answer that surprises you.

- Process 1: _____ → Watch / Murmuration — because

- Process 2: _____ → Watch / Murmuration — because

- Process 3: _____ → Watch / Murmuration — because

- Process 4: _____ → Watch / Murmuration — because

- Process 5: _____ → Watch / Murmuration — because

The one you most want to "add AI" to: _____. Is it a watch you're trying to make flock, or a murmuration you've trapped in a watch?

Tool 2 — The Loop Map Worksheet

Pick one real workflow and answer all five questions about it. Together they are a map of your system that you made yourself.

Q1 — Stock or flow. The one number your team watches most:

_____. Is it a level (stock) or a rate (flow)? _____. The stock it may secretly be draining: _____

Q2 — The delay. Your biggest lever: _____. How long after you pull it before you see the effect? _____. (Don't pull it again inside that window.)

Q3 — Leverage. The three places you're pushing hardest: (1) _____ (2) _____ (3) _____. Which is a *rule or incentive* (high leverage) versus a *number or more effort* (low leverage)? _____. Move your effort to the rule.

Q4 — Purpose (POSIWID). Watch what your AI actually does all day. The one-sentence purpose its *behavior* reveals: _____. Set beside your mission statement, the gap is: _____

Q5 — Conway's Law. Sketch your org chart in sixty seconds. The one team boundary your agents must cross and probably won't, because the humans don't: _____ . That's your first fix.

Tool 3 — The Orchestrator's Audit Score Sheet

Score each Stage statement from Appendix A: 0 = not yet, 1 = in places, 2 = consistently true.

- Statements 1-10 scores: _ _ _ _ _
- My Stage (first gate I can't clear): _____ (Dabbling / Assisting / Orchestrating / Operating / Native)

Score each Layer's two statements (0-2) and total each out of 4:

- Systems thinking: _ + _ = ____
- Context: _ + _ = ____
- Guardrails: _ + _ = ____
- Knowledge loops: _ + _ = ____
- Automation: _ + _ = ____
- My weakest layer (lowest total; ties go to the lower layer): _____

My result, in one sentence: "We are a Stage-____ organization whose weakest layer is _____ — so our next move is _____."

Tool 4 — The Orchestration Stack Go / No-Go Checklist

Before granting an agent more autonomy, check each layer. The weakest unchecked box is your ceiling — fix it before you grant more.

- ☐ Systems thinking — we mapped the workflow and decided where good decisions get made.
- ☐ Context — the agent can reach curated, current, business-specific knowledge, and the data is fit to feed it.
- ☐ Guardrails — hard limits are enforced in code: budget ceiling, allowed actions, kill-switch, human checkpoint for risky moves.
- ☐ Knowledge loops — outcomes feed back so behavior improves instead of repeating.
- ☐ Automation — only repeatable steps are automated; judgment calls are reserved for humans; the process was redesigned before it was sped up.

Lowest unchecked layer (the ceiling): _____

Tool 5 — The Orchestrator's Scorecard

Pick one AI initiative and put a real number in every box. If you can't compute one, that gap is the finding.

Initiative: _____

- Outcome-per-dollar (cost per *successful* outcome, retries included, vs. the human baseline): _____
- Autonomy ratio (share that runs without a human): _____
- Intervention rate (how often a human must step in): _____
- Time-to-trust (how long before you'll let it run with less oversight): _____
- Blast-radius-contained (how far damage spreads when it goes wrong):

The one box that's reddest — and what you'll do about it this quarter:

Tool 6 — The 30-Day Orchestration Sprint Plan

Pick one workflow worth the month. High pain, high measurability. Then run the six moves.

Workflow chosen: _____

- ☐ **Days 1-5 • Map.** Draw the workflow end-to-end on a wall. Mark every handoff, wait, and decision. (Day 1 is mapping — nothing else.) Done: ____
- ☐ **Days 6-10 • Add AI.** Put a model where it removes the most drudgery, not where it's flashiest. Done: ____
- ☐ **Days 11-15 • Build context.** Give it the business knowledge it needs; fix the data that isn't fit to feed it. Done: ____
- ☐ **Days 16-20 • Set guardrails.** Budget cap, allowed actions, validation gate, kill-switch, human checkpoint. Done: ____
- ☐ **Days 21-25 • Wire the loop.** Make this week's outcome improve next week's behavior. Done: ____

- ☐ **Days 26-30 • Measure.** Fill the five-metric scorecard with real numbers; compare to the baseline; decide keep / fix / kill. Done: ____

What worked, what you'd template for the next team: _____

Tool 7 — The 35 Put-It-to-Work Actions

The whole book, as a checklist. Tick one off each week and you have eight months of compounding moves.

Part I — The Inversion

- ☐ Ch 1 — Map one core process onto the Watch ↔ Murmuration spectrum; name the single speck of dust that would seize it.
- ☐ Ch 2 — List your five most critical processes; tag each watch or murmuration; defend any that surprise you.
- ☐ Ch 3 — List three things you still treat as "expensive to build" that are now cheap; reprice one this week.
- ☐ Ch 4 — Find one place AI made you busier but not better; write down *why* the speed didn't convert.
- ☐ Ch 5 — Write your one-line answer: when intelligence is free, what stays scarce for us, and what discipline protects it?

Part II — The Trap

- ☐ Ch 6 — Take a recent "win"; trace one second-order effect downstream; name what it cost elsewhere.
- ☐ Ch 7 — For one AI initiative, write the four constraints it must respect *before* it optimizes anything.
- ☐ Ch 8 — Give every autonomous agent a kill-switch and a ceiling; audit your most active one today.
- ☐ Ch 9 — Pick one agent workflow; compute cost per *successful* outcome (retries included) against the human baseline.
- ☐ Ch 10 — For your most autonomous agent, list what it can touch — tools, data, money — and cut the blast radius this week.
- ☐ Ch 11 — Run a one-week amnesty to map the AI people *actually* use; then pick the one workflow to orchestrate first.

Part III — The Shift in Seeing

- ☐ Ch 12 — Your most-watched number: stock or flow? Name the stock it's secretly draining.
- ☐ Ch 13 — Your biggest lever: how long until you see the effect? Mark the lag; don't pull again inside it.
- ☐ Ch 14 — Your three hardest pushes: which is a rule/incentive (high leverage) versus a number/effort (low)? Move effort to the rule.
- ☐ Ch 15 — Watch what your AI does all day; write the purpose its *behavior* reveals; name the gap with your mission.
- ☐ Ch 16 — Sketch your org chart in sixty seconds; circle the one boundary your agents must cross and won't.

Part IV — The Orchestration Stack

- ☐ Ch 17 — Score your last AI project layer by layer; the weakest layer is your autonomy ceiling — fix it first.
- ☐ Ch 18 — Find one decision you personally bottleneck; design the system that makes it without you.
- ☐ Ch 19 — Write the ten things an agent must know about your business; check whether that data is actually accessible.
- ☐ Ch 20 — For your most autonomous agent, list three things it must never do; enforce them in code this week.
- ☐ Ch 21 — Pick one agent; add one loop so this week's outcome improves next week's behavior.
- ☐ Ch 22 — Before automating anything, ask: would I be proud if this ran 1,000× faster as-is? If not, fix it first.
- ☐ Ch 23 — Redefine one role from "produces output" to "designs the system that produces output."
- ☐ Ch 24 — Audit where the system can make a final, unrecoverable call on your behalf; for each, ask "if this is wrong, can we take it back?" Leave the reversible ones with the machine; pull every one-way door back to a named human.
- ☐ Ch 25 — Count the decisions on your desk this week; pick three the *system* should own by next quarter.
- ☐ Ch 26 — Name one role being hollowed out; draft the 90-day path that

moves that person up the stack.

Part V — The AI-Native Organization

- ☐ Ch 27 — Locate your org on the staircase honestly; name the one thing blocking the next step.
- ☐ Ch 28 — Find two AI tools that should talk and don't; connect them this month.
- ☐ Ch 29 — Write the one-paragraph "day in the life" of your best team twelve months out; reverse-engineer step one.
- ☐ Ch 30 — Build the five-metric scorecard for one initiative; put a real number in each box; show it to your own Diane.
- ☐ Ch 31 — Turn one manual approval into an automated gate embedded in the workflow itself.
- ☐ Ch 32 — For one AI decision a customer feels, add a way to see why — and to appeal it.

Part VI — The Last Premium Skill

- ☐ Ch 33 — Name the three judgment calls in your org that must never be delegated; defend why, in writing.
- ☐ Ch 34 — Pick the one workflow; start Monday; Day 1 is mapping, nothing else.
- ☐ Ch 35 — Write your one orchestration principle on one card. That's the system you'll lead by.