THE BUILDER'S STACK · THE FRONTIER · ORCHESTRATION

The Orchestration Plan

One page that turns a big task into a fleet you can run: the cut-list, the shape, the verification stages, the budget, and the gates where you stay in command. Fill it before you spawn anything; keep it with the run notes afterward. The second fleet starts from the first plan.

PROJECT / TASK

RUN BY

DATE

1 The task and its success bar

Name the deliverable and the bar it must clear. If you cannot state the bar, the fleet cannot hit it.

THE TASK, IN TWO SENTENCES

DONE MEANS

2 The cut-list

Five to eight items, each with a contract. If two items share a must-not-touch entry, re-cut until none do.

ITEM 1

WHAT THIS AGENT PRODUCES

IT GETS

IT RETURNS

IT MUST NOT TOUCH

ITEM 2

WHAT THIS AGENT PRODUCES

IT GETS

IT RETURNS

IT MUST NOT TOUCH

ITEM 3

WHAT THIS AGENT PRODUCES

IT GETS

IT RETURNS

IT MUST NOT TOUCH

ITEM 4

WHAT THIS AGENT PRODUCES

IT GETS

IT RETURNS

IT MUST NOT TOUCH

ITEM 5

WHAT THIS AGENT PRODUCES

IT GETS

IT RETURNS

IT MUST NOT TOUCH

ITEM 6

WHAT THIS AGENT PRODUCES

IT GETS

IT RETURNS

IT MUST NOT TOUCH

3 The shape

Pick the pattern the work actually has, not the cleverest one.

- ☐ **Fan-out** — independent items run in parallel, gathered at the end
- ☐ **Pipeline** — each item flows through stages, no barriers between stages
- ☐ **Writer-verifier** — every item gets a checker with authority to fix

- [ ] **Judge panel** — several independent attempts or judgments, then a vote or synthesis
- [ ] **Orchestrator-worker** — a lead agent decomposes, spawns workers, and synthesizes

WHY THIS SHAPE FITS THIS WORK

4 Verification

Deterministic checks first: scripts, greps, and counts are cheaper than judgment and never get bored. Spend model judgment only where it is genuinely required.

DETERMINISTIC CHECKS (SCRIPTS, GREPS, COUNTS)

JUDGMENT CHECKS (LENSES, REFUTERS, WHAT EACH LOOKS FOR)

- [ ] Verifiers have authority to fix in place (not just report)

5 Budget and tiers

Parallel buys time, never tokens. Cap discovery, never verification, and put the strongest model where judgment concentrates.

HARD CAP FOR THIS RUN

AT THE CEILING, WE

MODEL TIERS BY ROLE (MECHANICAL ROLES VS. JUDGMENT ROLES)

6 Gates and abort

Two gates always exist: before anything irreversible, and before anything expensive. The one-glance surface answers 'is it going well' without reading transcripts.

GATE: BEFORE ANYTHING IRREVERSIBLE

GATE: BEFORE ANYTHING EXPENSIVE

THE ONE-GLANCE SURFACE SHOWS

WE ABORT THE RUN IF

7 The resume story

Assume the run gets interrupted. What survives, what replays instantly, and how do you prove nothing was lost or duplicated?

IF THIS RUN DIES MIDWAY

8 After the run

The plan is a living template. Write down what it got wrong while the run is fresh.

RUN DATE

THE ONE THING THE PLAN GOT WRONG

KEEP FOR THE NEXT FLEET