

ADW Kickoff: 2-Hour Facilitator Run-sheet

For whoever runs a team's first ADW session. The whole plan fits on one page so you can run the session straight off it. Prefer paper? Print this page, or grab the [one-page PDF](#).

The one rule for the room: everyone leaves having run the loop themselves on their own code. A demo they watch is not the goal; hands on keyboards is.

Two messages that repeat all session: AI does the building, people approve every change (human-gated), and it runs on a flat, unmetered subscription before any metered spend (local-first).



COMPANION MATERIALS

Present from the [slide deck](#), hand out the [PDF](#), and point the team at the [Executive Briefing](#) for the framing.



THE HANDS-ON RUNS ON A SHARED PLAYGROUND, NOT PERSONAL REPOS

A brand-new repo has no `lat.md/`, no ADW commands, and no `lat` key, and onboarding a repo is itself later work (ADW Track Module 7). So the hands-on runs on the **ADW playground**, a small pre-grounded repo everyone clones:

```
git clone https://dev.azure.com/wasteology/Wasteology/_git/adw-playground
```

Onboarding each person's own repo comes later, in the weekly cadence, with the champion and the `project-onboarder` agent.

The agenda

Time	Segment	What you do
0:00-0:10	Why we're here	Five-minute frame: AI builds, humans approve, local-first. Set the promise out loud: by the end everyone will have run this on their own repo. Do not oversell.
0:10-0:30	Live demo: one full loop	Take a small, real task on a real repo. Kick the loop and narrate lightly as it goes classify, plan, implement, check, commit, draft PR. Then review and merge the draft PR in front of them. This is the payoff, so lead with it.
0:30-0:50	The spec is the leverage	Now open the markdown spec: "here is where that came from." Show a weak spec next to a strong one for two minutes. Message: garbage spec in, garbage PR out. The spec is the skill.
0:50-1:05	Grounding with lat.md	Open the playground's <code>lat.md/</code> . You demo <code>lat search</code> (it needs an API key); everyone can run <code>lat locate</code> and <code>lat check</code> themselves (no key). The point, not the syntax: the agent is not guessing, and a broken reference cannot merge.
1:05-1:15	Break	Ten minutes.
1:15-1:50	Hands-on (the heart of it)	Everyone opens Claude Code on the cloned ADW playground . Minimum win for each person: run <code>lat locate</code> and <code>lat check</code> , then ask Claude a grounded question and watch it answer from <code>lat.md/</code> . Stretch, if their machine is set up: run <code>/implement specs/task-2-fix-rounding.md</code> for a full loop, or <code>/feature <idea></code> then <code>/implement</code> . Circulate and pair.

Time	Segment	What you do
1:50-2:00	Where this goes, and Monday	Three-minute teaser of the "one terminal, many projects" hub (that is Phase 2, not a tool for tomorrow). Then next steps: the weekly cadence, the Academy for self-serve, which 2 to 3 lighthouse repos we start on, and who the champions are. Close on the two rules and the maturity ladder.

Say this at the three anchor moments

- **After the demo:** "Every line an agent wrote lands as a draft PR that a person reviews and merges. Nothing ships on its own."
- **On the spec:** "The agent is only as good as the spec. Most of the skill we are learning is writing the spec, not prompting."
- **On lat.md:** "This is how the agent reads our intent instead of guessing, and the gate is why our docs cannot quietly rot."

Keep it light (resist the urge to go deep)

- The `ops` CLI and the two cloud execution models are Phase 2. One teaser at the end, not a tour.
- The `lat` CLI is three commands today, shown in service of the loop.
- Cost and governance get one sentence: every change is a human-merged draft PR, and it runs on your flat subscription. The detail lives in the Academy.

Before the session (do these the day before)

- ☐ Everyone has **Claude Code installed and signed in** to their subscription. This is the number one thing that eats live time, so send setup instructions the day before.
- ☐ Everyone has the **lat CLI installed** (pinned version, per the setup message) and has **cloned the playground**.
- ☐ **(Facilitator only, optional) Confirm lat search works on your machine** so you can demo semantic search. It needs a Gemini API key (an `AIza...` key from Google AI Studio, `aistudio.google.com/app/apikey`) *and* the Gemini-patched `lat` (the pinned install). If you already run `lat search` in wg-orchestration you are set. The team does not need this at all; they use `lat locate` and `lat check`, which need no key.
- ☐ Your **demo task is chosen** and you have done a dry run so the live loop does not surprise you.
- ☐ A **recorded fallback** of the loop is ready in case the live run stalls or the network drops.
- ☐ The **2 to 3 lighthouse repos are pre-named**, so "what's next" is not a debate in the room.
- ☐ The **companion tabs are open**: the deck, the Executive Briefing, and Academy Modules 0, 1, 3, 4.

After the session

- Send the follow-up path: Academy **Module 0** (why), **Module 1** (Claude Code locally), **Module 3** (spec-first), **Module 4** (the local loop).
- Confirm the weekly cadence: one module or topic plus live office hours on the team's real code.
- Name the champion, and agree the first lighthouse repo each person will onboard.
- Schedule the **"onboard your own repo"** session (ADW Track Module 7): scaffold `lat.md/` + the lat-check gate + the commands on a real repo, with the champion and the `project-onboarder` agent. This is the step that turns the playground practice into ADW on real work, and it is deliberately not day one.

This session is the compressed, demo-first version of the ADW Track arc (why, local tooling, grounding, spec-first, the loop). The two hours are the on-ramp; the Academy is the follow-through.