

Jordan Reyes

Ops/Partnerships · 4 weeks · 3x per week, 45 min per session

Name	Jordan Reyes
Track	Ops/Partnerships
Start	2026-06-13
Duration	4 weeks
Sessions	3 per week
Commitment	45 min per session

Lighthouse Chart — Jordan Reyes, Ops Lead

1. Opening / Bearings

You are spending real hours every week on work that follows the same shape every time: find the right words for a partner you have never met, pull scattered notes into a coherent status update, write down a process that lives in your head so someone else can follow it. The goal here is not to hand that work to a machine — it is to build a small, reliable set of AI-assisted workflows that cut the repetitive load so you can spend your judgment on the parts that actually need it.

Four weeks. Three sessions a week. By the end, you will have sent five AI-drafted partner outreach emails, a weekly summary process that takes under 20 minutes, and at least one SOP in active team use with an AI-drafting step baked in. Those are the three outcomes this plan is built around, in that order.

You already know what a prompt is, you have used ChatGPT occasionally, and you work in Notion and HubSpot daily. That is enough to start. The gap between where you are and where this plan ends is not technical — it is about building the habit of giving AI enough context to do useful work, and then storing that context so you do not start from scratch every session.

That is what the next four weeks are for.

2. Track Context

Your work sits squarely on the Ops and Partnerships track. The defining feature of this track is that the highest-value AI use is not creative generation — it is systematic, repeatable output from structured inputs. You are not trying to write something original every time; you are trying to produce something consistently good from a known set of inputs (a partner name, a mutual benefit angle, a week's worth of raw notes) without burning 45 minutes doing it.

The two failure modes on this track are worth naming early so you can avoid them.

The first is the blank-page prompt. You open Claude, type "write me an outreach email," get something generic, decide AI is not useful for this, and go back to writing it yourself. The fix is specificity: the more context you give the model about the partner, the angle, and the tone, the more the output looks like something you would actually send. The ready-to-paste prompts in this chart are built to give you that specificity from day one.


The second is the one-session wonder. You build a great prompt, use it once, and then cannot find it the next week. The fix is storage: every prompt you refine goes into a Notion page immediately, and by week two it also lives in OpenBrain so it travels with you across Claude sessions. That is the whole system — good prompt, stored prompt, reused prompt.

The upcoming partner onboarding project is a natural extension of this track. It is not in scope for the four-week plan (you listed building a full automation platform as a non-goal, and that is the right call), but the stage-mapping work in week three gives you a foundation to build on after this plan ends.

One more thing: you have HubSpot in your stack. Claude cannot write directly into HubSpot, but the workflow is fast — draft in Claude, paste into HubSpot, send. That two-step process is the right shape for now. It keeps a human in the loop on every send, which matters for partner relationships, and it does not require any IT approval beyond what you already have.

The rest of this plan is redacted.

You are seeing the first two pages. The full plan is yours when you create one.



Create yours at lighthouse.dom.net