

A BARELY SHIPPING FIELD GUIDE


The Three Levels Field Guide

A prompt is a task.

A workflow is a process.

An agent is a system that acts.

Run any process in your go-to-market engine through this guide and know exactly what to build.



Nathan Thompson


Barely Shipping | Systems-Led Growth

Companion to the podcast episode: Chat vs. Workflows vs. Agents

REFERENCE. TAPE THIS TO THE WALL.

The ladder

Three levels. Each one earns its place, and each one gets misused when it does the next level's job.



AGENT *A system that acts.*

It observes a situation, evaluates it, and chooses between paths.

- Inbound lead routing
- Support resolve / refund / escalate
- Bug and request triage

Agents scale decisions.

WORKFLOW *A process, codified.*

An expert ran it manually, start to finish. Now anyone on the team gets the expert's output.

- Monday content brief batch
- Call transcript to follow-up email
- Draft QA against brand rules

Workflows scale a process.

CHAT *A task, done faster.*

Individual, non-repeatable work. You drive, the model assists.

- Answer a specific email
- Punch up an intro
- Research where something lives

Chat scales a person.

**You can't skip levels.
And you shouldn't use a higher level where a lower one belongs.**

Defined or decides

Run any process through two questions before you build anything.



THE WARNING LABEL

If you're not sure whether your process is defined or decides, you're not ready to build either one yet. Run it manually until you are. Building an agent for a defined process means paying the agent tax (unpredictability, compounding errors, harder QA) for zero benefit.

WORKSHEET. PRINT IT, WRITE ON IT.

Write the SOP

Pick one process your team repeats every week. Write the SOP below, step by step. Every time you catch yourself writing "it depends," "usually," or "review it first," check the box on that row.

Process:	It depends?
1. _____	☐
2. _____	☐
3. _____	☐
4. _____	☐
5. _____	☐
6. _____	☐
7. _____	☐
8. _____	☐
9. _____	☐
10. _____	☐
11. _____	☐
12. _____	☐

Score it

0 boxes, and you finished: You just wrote a workflow spec. You're one build away from getting that time back.

1 to 3 boxes: Workflow with human checkpoints. Mark exactly where a person steps in.

Boxes everywhere, or you couldn't finish: The judgment is the process. That's human territory today (and maybe an agent much later).

Where each level lives

	CHAT	WORKFLOW	AGENT
CONTENT	Punch up an intro. Fix a paragraph. Find three sources.	The Monday batch: SERP research, library links, gap analysis, template. Ten briefs in ten minutes, then human editing.	QA-and-route: evaluates each draft and decides how much human attention it needs (publish, light edit, or rewrite).
SALES	Phrase a reply to one tough prospect email.	Transcript in, follow-up out: extract pain points, map to value props, generate the email and one-pager.	Inbound routing: enrich, score, then decide the path. Rep with prep brief, auto-response, or nurture.
SUPPORT	Summarize one angry ticket before you reply.	KB answer bot: receive question, search knowledge base, compose answer. Dynamic inputs, defined process.	Full resolution: decides whether to resolve, refund within policy, escalate to a human, or open a bug ticket.
PRODUCT	Draft release notes from a changelog.	Feedback digest: collect the week's requests, cluster, summarize for the team.	Triage: evaluates severity, checks duplicates, decides the route (hotfix queue, backlog, or known-issue reply).

THE ONE EVERYBODY GETS WRONG

A bot that searches your knowledge base is a workflow with dynamic inputs. The questions vary. The process doesn't (receive, search, compose). Call it an agent and you'll build agent-grade complexity for workflow-grade work.

THE HONEST CONS. EVERY LEVEL BREAKS SOMEWHERE.

Where each level breaks

CHAT

- ✗ Outsourced thinking. The blank window invites you to hand over strategy, not just production. That's where slop comes from.
- ✗ Team silos. Ten people, ten prompts, zero visibility into what produced a great (or terrible) output.
- ✗ False context confidence. The answer looks complete. The input wasn't.
- ✗ Scales the person, not the team. The process lives in one head.

WORKFLOW

- ✗ The expert bottleneck. The person who knows the process best often can't build, and the builder doesn't know the process. Expert and builder are different jobs.
- ✗ Brittleness. Workflows codify a process at a point in time, and processes go stale. Outputs still look consistent while quality quietly degrades. Consistency is the feature and the trap.

AGENT

- ✗ Outsourced judgment, with autonomy attached. The agent produces something good while missing gaps an expert would catch. Review late and you scale the blind spot.
- ✗ Compounding error (see below).
- ✗ Silos, worse. Personal agents with personal context give you chat's silo problem plus decision-making authority.

THE COMPOUNDING ERROR MATH

$$0.95 \times 0.95 \times 0.95 \times 0.95 \times 0.95 \times 0.95 = 0.74$$

A 95% accurate decision made six times in a row is 74% accurate end to end. A workflow makes one pass and surfaces one error. An agent chains decisions, and error rates multiply across the chain.

That's how agents fail. Not one big mistake. A chain of reasonable calls that lands somewhere you'd never have gone.

THE AGENT ZONE. WHERE 'DECIDES' BEATS 'DEFINED.'

Should this be an agent?

	PREDICTABLE INPUTS	UNPREDICTABLE INPUTS
STABLE GOAL	WORKFLOW TERRITORY Write the SOP. Build it once.	THE AGENT ZONE • Inbound leads • Support tickets • Comment triage • Account signals
UNSTABLE GOAL	HUMAN TERRITORY	HUMAN TERRITORY If the goal keeps moving, no system saves you.

Agents belong where the inputs are unpredictable but the goal is stable.

THE READINESS CHECKLIST (ALL FIVE MUST BE TRUE)

- ☐ Inputs are unpredictable. You can't script what arrives.
- ☐ The goal is stable. You know exactly what a good outcome looks like.
- ☐ You've run this process manually enough to judge its outputs on sight.
- ☐ Errors are catchable before they compound. A review point exists.
- ☐ A workflow genuinely can't do it. You tried to write the SOP and failed.

WHAT'S NEXT

Go deeper

The book

Pipes Before Chocolate walks the full build order: the four audits, the minimum viable factory, and the 30-day plan, week by week.

The podcast

Barely Shipping. Next episode goes deep on defined vs. decides, the single highest-leverage architecture decision you'll make.

The newsletter

One build, one honest admission, one thing that broke. Weekly.

Before you close this

Go back to page 4. Pick the process. Write the SOP. Circle every "it depends." That exercise costs nothing and tells you exactly what to build first.

Obvious disclosure: I build workflows and systems for a living, so weigh my bias toward level two accordingly. The reasoning is on these pages, not in my invoice.

Nathan Thompson

Barely Shipping | Systems-Led Growth