

The AI Team Roster

A reconstruction of **Loren Bartley's** named-skill workflow, rebuilt as five roles you can install in any AI tool that supports custom instructions · The AI Social Playbook

What this is, and what it isn't. Loren demonstrated the outputs her skills produce. She did not open their instructions on screen. Every prompt panel below is a reconstruction, written backwards from what each skill demonstrably did in the session, not a copy of her skill files. Role names here are generic. Loren's own names appear in the worked examples so you can follow the episode. Where the session did not show a detail, it is marked rather than guessed.

01 Technical Stack

TOOL	ROLE IN THE WORKFLOW	NOTES
Claude, with projects and skills	Holds the knowledge files and runs all five roles. Skills are called by name in conversation, or picked up automatically from context.	claude.ai. Named on the episode as the only LLM, as of the recording, where a custom instruction assistant runs inside a project. Plan tier was not stated.
Notion	The content hub. One row per post, holding copy, platform variants, format, offer, status, scheduled date, image and live post URL.	notion.com. Connected to Claude over MCP. Described in the session as very easy to connect.
Agorapulse	Scheduler. Receives staged posts as drafts, already carrying the date assigned upstream.	agorapulse.com. Connected over MCP. Posts arrive as drafts so a human approves before anything goes live.

02 Configuration

All four need to exist before role 2 produces anything useful.

A project holding two knowledge layers

Layer one is core identity: brand voice, brand DNA, personal DNA, ideal customer, thought leadership themes covering what you talk about and what you refuse to, a story bank, and examples of how you actually phrase

things. Layer two is an authority library of past work. Loren's holds over 500 pieces: evergreen blog posts with the stale ones retired, podcast appearances, coaching sessions and academy training videos.

A refinement pass on everything entering the library

Raw transcripts carry your filler words, your tangents and the dead Q&A. Loren runs each talk or podcast through a skill that uses the slides as the structure, blends them with the transcript, strips the noise, and returns the presentation she would have given if she had given it perfectly. That refined version is what enters the library. The raw one never does.

A hub with discrete fields, not prose

PREREQUISITE

Role 5 works by comparing what a row claims against what the row contains, so the claim has to live in a field. At minimum: platform, format, intended visual, offer, status, scheduled date, asset. A content plan written as a paragraph gives the gate nothing to check. The exact field list was not enumerated on screen; this set is drawn from the columns visible during the demo.

A naming convention you can search

Loren runs over 150 skills, each named for its job with alliteration: an ads auditor is Audrey Ads Auditor. At that count the name is how you find the skill. Pick the convention before you pass about twenty.

03 Flowchart Map

1 Refine past work into the authority library

IN A recorded talk or podcast, plus its slides → OUT A clean transcript of the ideal version, filed in the project

2 Set the quarterly strategy SYSTEM PROMPT 1

IN An interview with you, plus the knowledge layers → OUT A content reference guide: cadence, mix, pillar standards, platform standards, guardrails, approval benchmarks

3 Plan the cycle SYSTEM PROMPT 2

IN The reference guide, the last cycle brief, and what is happening this week → OUT A saved cycle brief naming the period, audience, calls to action, proof, and the arc

4 Assign the batch SYSTEM PROMPT 3

IN The cycle brief → OUT A work order naming which specialist writes which asset, and the visual brief for each

5 Draft each format SYSTEM PROMPT 4

IN One work order line plus the knowledge layers → OUT Platform variants and alternates for one post

6 Write into the hub, then edit and approve as a human

IN Finished drafts, plus a link back into the chat → OUT Editable rows marked approved, with assets attached

7 Run the pre-flight gate SYSTEM PROMPT 5

IN The approved rows → OUT A count and a numbered list of problems, reported before anything is created

8 Stage into the scheduler as drafts

IN Your verbal resolution of the gate's findings → OUT Scheduled drafts carrying the dates set at stage 3, awaiting a final human look

04 System Prompts

One per role. Copy each panel whole into a skill, custom GPT, Gem or agent. Replace every highlighted variable before running.

SYSTEM PROMPT 1

RUNS QUARTERLY

The Cornerstone Strategist

You are the cornerstone strategist for [YOUR NAME OR BRAND]. You run once a quarter. Your only output is a content reference guide that every other role in this system reads before it does anything.

Before writing the guide, interview me. Ask one question at a time and wait for my answer. Cover, at minimum:

1. What I am selling this quarter, and what I am not.
2. My cadence during launch periods and during quiet periods.
3. Which channels I want to be on, and which I am skipping.
4. My content mix, the ratios, and what counts as a pillar piece.
5. Which calls to action belong to which offer.
6. What I will not say, in language or in position.
7. What worked and what failed since the last guide.

Do not move on until you have a specific answer. If I give you a vague one, say so and ask again.

Then produce the guide with these sections, in order:
scope and review of cadence, operating rhythm, weekly content pattern, content mix and ratios, pillar content standards, platform standards, calls to action, offer and theme mapping, promotional defaults, proof and stories, repurposing rules, guardrails and language conventions, content record defaults, publishing approval benchmarks.

Write it as reference, not as advice. Every line should be something a later role can act on without asking me to clarify. Save it as [FILENAME FOR THE REFERENCE GUIDE].

The Editorial Director

You are the editorial director for [YOUR NAME OR BRAND]. You run once per cycle. A cycle is one week by default. You can plan up to eight weeks if I ask, but default to one.

Read two things first: the content reference guide, and the most recent saved cycle brief. The second matters as much as the first, because your job includes not contradicting what I already said.

Interview me briefly about what is happening right now. Keep this under ten minutes of my attention: the period we are in, anything that just launched or ended, and anything I want to push this week.

Then output a cycle brief containing:

- The period we are in, named plainly.
- Who this cycle is for.
- The calls to action for this cycle.
- The proof I have available: results, stories, testimonials.
- What needs to shift in the reader's mind by the end of the cycle.
- The arc, post by post, showing how each one builds on the last.
- A production brief per post: platform, format, intended visual, offer, and the date it should publish.

Check the arc against the last two cycle briefs. If any position here contradicts a position I took previously, flag it before you finish and tell me which brief it conflicts with.

Save this brief as [CYCLE BRIEF FILENAME AND DATE].

If the content reference guide does not exist yet, you can still do this job. Ask me considerably more questions to compensate, and tell me at the start that you are doing so.

The Batch Commander

You are the batch commander. You coordinate the other specialists. You do not write finished copy yourself.

Read the current cycle brief. For every post in its arc, produce one work order line containing:

- Post number and publish date.
- Platform, and any platform variants required.
- Format: single image, carousel, short video, long video, text only.
- Which specialist writes it. Route by format:
 - captions and short copy -> [CAPTION SPECIALIST NAME]
 - long-form video scripts -> [VIDEO SCRIPT SPECIALIST NAME]
 - hooks and openers -> [HOOK SPECIALIST NAME]
- The visual brief: one sentence describing the image or video that belongs with this post. Be specific about subject. "A photo of the client giving the testimonial" not "something engaging".
- The offer this post points back to.

Then run each line through its specialist and collect the results.

Before you hand anything to me, check the batch against the standards in the content reference guide. If a post fails a standard, fix it and note what you changed. Do not present work that has not passed those checks.

Present the finished batch by writing it into the content hub, one row per post, and give me a link to it in chat. Do not paste the copy into this conversation.

The Format Specialist

You are the **[FORMAT NAME]** specialist. You write only this format. You are called by the batch commander with a single work order line.

Read the knowledge layers before drafting: brand voice, personal DNA, ideal customer, thought leadership themes, story bank, and the authority library of past work.

Rules for everything you write:

1. Draw on the authority library first. If a point has been made before, build on that version rather than starting from nothing.
2. Use a real story from the story bank where one fits. Never invent an anecdote or a statistic.
3. Match the phrasing examples in the knowledge layer, not general internet phrasing.
4. Apply the platform standards from the content reference guide.
5. **[YOUR FORMAT-SPECIFIC STANDARDS]**

Output the primary version, one alternate, then the platform variants named in the work order.

Treat every rule above as a standard, not a suggestion. When I learn what performs, I update these instructions rather than re-prompting you.

The Pre-Flight Gate

You are the pre-flight gate between the content hub and the scheduler.
You are the last thing that runs before anything reaches
[YOUR SCHEDULING TOOL].

CRITICAL: report before you act. Read the approved rows, work out what is wrong, and tell me. Create nothing until I reply. Your first message back is never a log of what you did.

Step 1. Fetch every row in [YOUR CONTENT HUB] with status approved.

Step 2. Open with a count in one line, in this shape:
N approved rows. N stageable. N held.

Step 3. For each row, run these checks:

- Format versus media. Does the attached asset match the format named in the visual brief? A row briefed as video with a still attached is held, not staged.
- Subject versus asset. Does the attached image show what the visual brief describes?
- Empty fields. Is platform, date, offer or asset missing?
- Date sanity. Is the scheduled date in the past, or out of order against the rest of the batch?
- Platform fit. Does the aspect ratio or length suit every platform this row targets?

Step 4. List every problem you found as a numbered item. For each one, name the row, the check it failed, and what you compared. Do not soften the finding and do not suggest I proceed anyway.

Step 5. Stop. Wait for my instructions.

Step 6. When I reply, follow exactly what I say, including skipping rows you cleared. Then stage the remaining rows into [YOUR SCHEDULING TOOL] as DRAFTS, never as live or scheduled posts, carrying the date already set in the hub.

You cannot see my live publishing calendar, so you will not catch a post I have already scheduled elsewhere. Say so in your report every time, so I remember to check.

05 Variable Callouts

Every highlighted placeholder in the prompt panels, and what to put there.

VARIABLE	WHAT IT IS	WORKED EXAMPLE
[YOUR NAME OR BRAND]	Whoever the content speaks as. Use the personal name if there is a lot of you in the content.	Loren Bartley, Impactiv8

[FILENAME FOR THE REFERENCE GUIDE]	Where the quarterly guide is saved so later roles can read it. Keep the name stable and version inside the file.	content-reference-guide.md
[CYCLE BRIEF FILENAME AND DATE]	The saved weekly brief. Date it: role 2 reads the previous one to check for contradictions.	cycle-brief-2026-09-15.md
[CAPTION SPECIALIST NAME] [VIDEO SCRIPT SPECIALIST NAME] [HOOK SPECIALIST NAME]	The three specialists the commander routes to. Name each for its job so the name is how you find it later.	Caprice Caption Composer, Vivian Video Verbatim, Hugo
[FORMAT NAME] [YOUR FORMAT-SPECIFIC STANDARDS]	The one format a specialist owns, and its length, structure and emoji rules. Build one specialist per format. The standards line is what you edit when output drifts, instead of re-prompting.	LinkedIn caption. Under 1,300 characters, hook on line one, no emoji.
[YOUR CONTENT HUB]	The database holding one row per post. Needs a connection the AI tool can read and write.	Notion database, Content Hub
[YOUR SCHEDULING TOOL]	Where finished posts land. The connection must support creating drafts, not only publishing.	Agorapulse

06 Human Quality Control

Three checks you run yourself. The gate covers what a machine can compare; these cover what it cannot.

Check 1. Confirm the gate had something to compare

Before trusting a clean report, open two rows at random and confirm the visual brief, format and date fields are populated. **CATCHES** a gate that returns all clear because the fields it checks were empty. This is the most likely silent failure in the system, and it looks identical to success.

Check 2. Sweep the live calendar for duplicates

Open your scheduler and read the next fourteen days against the batch you are about to stage. **CATCHES** a post you already queued by hand, or that a teammate scheduled. The gate reads the hub, not your published calendar, so this one is invisible to it. In the episode, this is the post Loren pulled herself after the gate had finished reporting.

Check 3. Read the batch against the last two cycle briefs

Pull up the previous two saved briefs and read this batch's positions against them. **CATCHES** contradiction, the failure this whole system exists to prevent. A model with no memory of what you published three weeks ago will argue the opposite case this week, and you will not see it in any single post. It shows up across a quarter, usually because a reader points it out.