ARCHITECTURE NOTES · RESOURCE

The five-step brief-to-ship loop

A workflow you redesign every Monday is not a workflow. The leverage is the schema; the schema is the boring part on purpose.

Companion to: The five-step brief-to-ship process for AI content cadence.

Why the schema is the leverage

The five-step loop below runs every week with the same shape, and the leverage is the schema rather than any single element of it. The director never walks into the week with an empty board. The drafter writes against a brief, not a vibe. The voice contract is enforced at every boundary instead of only at the end. The logs are structured around outcomes rather than triggers. Any of these turning ad-hoc breaks the loop, and the loop stops running itself.

- Step 1: The board is preloaded before the week starts
- Step 2: Strategy as a short scan
- Step 3: The drafter writes
- Step 4: Human review
- Step 5: Schedule, publish, monitor

The five steps in order

01 **Step 1: Board preloaded before the week starts**

On Monday morning the strategic queue already carries briefs in idea status, organised by specialist and pillar. The director never walks into the week with an empty board. The mechanism is a scheduled job that runs over the weekend, scans the cadence targets per specialist, and

02 **Step 2: Strategy as a short scan**

The director scans the queue, approves briefs that match this week direction, and rewrites the rest with a one-line note. Approved briefs route to the specialist who will draft them. Rewritten briefs return to the queue with the new framing. The pass is short because the brief shape is

03 **Step 3: The drafter writes**

A specialist agent (n8n + Claude in our stack) picks up an approved brief, reads the relevant context (audience, voice, prior work, sources), and produces a draft. The draft lands in a queue with platform, content, hook, and voice fields populated. A voice gate scans every draft on

04 **Step 4: Human review**

Review is keyboard-native. Up arrow, down arrow, approve, reject, request rewrite. Each draft shows the copy, the image preview, the source URL, and a voice-status badge. The job in this step is to decide what ships, not to write.

05 **Step 5: Schedule, publish, monitor**

Approval moves a draft to scheduled. A dispatcher cron checks regularly for due posts, runs one final voice scan as a publish-boundary check, calls the platform publishing layer, and marks the row published. The activity feed updates in real time. Failures retry; persistent failures

Loop you can run vs. workflow you redesign weekly

Same-shape loop (this is the leverage)

- Director starts from a draft to react against, not a blank canvas
- Voice gate enforced at three boundaries (entry, review, publish)
- Brief schema identical every week
- Drafts that fail the gate get rewritten before any human sees them
- Logs structured around outcomes (rows-written, voice-pass-rate)

Ad-hoc workflow (breaks each week)

- Empty board on Monday morning
- Voice judged only at the human review step
- Brief format negotiated case by case
- Operator spends review time on drafts the system already knows are broken
- Logs report API calls, not artifacts

What good looks like at each boundary

VOICE-PASS RATE AT GATE

Majority

Most drafts pass the entry scan; failures get rewritten before human review

DIRECTOR REVIEW TIME

Short scan

Strategy step is decisions, not form-filling, because the schema is fixed

PUBLISH-BOUNDARY SCAN

Final gate

One last voice check before the platform publishing layer is called

Operator runbook: install the loop in your team

- **Fix the brief schema first**
  Same fields every week: pillar, voice, hook, key points, CTA, source. The schema is the boring part on purpose. Fix it before anything else.
- **Build the weekend seed job**
  Scheduled job that scans cadence targets per specialist and seeds gaps with pillar-shaped templates. The director arrives Monday with a board to react against, not to fill.
- **Wire the voice gate at three places**
  Entry (drafter output), review (badge in the UI), and publish (final boundary check by the dispatcher). Three checks, same contract.
- **Make review keyboard-native**
  Up arrow, down arrow, approve, reject, request rewrite. The reviewer decides what ships; they do not write. The shape of the UI enforces the role.
- **Park repeatedly-failing drafts for manual editing**
  Drafts that fail repeated rewrites should not loop forever. Park them for a human to edit, and learn from the parked drafts what the brief needs to say differently.
- **Surface persistent failures with the exact error**
  When the dispatcher fails repeatedly, push the error string straight to the operator. Fast triage beats long log archeology every time.
- **Audit the loop weekly, not the drafts**
  The structural deliverable is the cadence the schema enforces. If the cadence holds, the drafts are working. If the cadence drifts, the schema needs work.

Where the loop breaks

01 **Brief schema goes ad-hoc**

The moment briefs start being negotiated case by case, the director loses the short-scan property and the loop becomes a meeting again.

02 **Voice gate only at the end**

If voice is judged only at human review, the operator burns time on drafts the system should have caught. Three boundaries exist for a reason.

03 **Empty Monday board**

Without the weekend seed job, the director starts from a blank canvas every week. Strategy time turns into form-filling time and the loop slips behind.

04 **Logs that report triggers**

If the dispatcher logs API calls instead of rows-written and publish-boundary outcomes, silent failures hide inside green status codes. The loop needs heartbeats, not triggers.

NEXT STEP

The schema is the leverage

Install the loop once and the cadence runs itself. The structural deliverable is the schema, beyond any individual draft.

Read the full architecture note ->