

TL;DR We run SEO like a content ops loop: pull GSC demand signals, validate what is winning the click in a real browser, confirm the freshest sources, then publish (or refresh) with one clear canonical page per intent. Measure again 7-14 days later.

1 Why we built this playbook

Most teams plan SEO content by vibes:

- "This topic feels hot."
- "A competitor wrote about it."
- "ChatGPT suggested it."

The predictable result is wasted output:

- You publish pages that do not match a query intent.
- You create multiple pages for the same intent and split ranking signals.
- You quote facts that were true 6 months ago but are now stale.

This playbook is how we reduce guesswork and turn SEO into a repeatable input for our wider AI production stack.

If you want the broader system view, see: [AI Content Ops System: From Brief to Measurement](#).

2 The three inputs we trust

Input A: Google Search Console (GSC)

GSC tells you what Google already tested you on:

- Which pages are earning impressions and clicks.
- Which queries you are showing for (even if you are not getting clicks yet).
- Which pages have a CTR problem even though they rank well.

Tooling note: we pull this data via GSC MCP (Search Console API). If you are doing it manually, export the same views from Search Console and keep the time window consistent across runs.

Input B: SERP snapshots (real browser)

GSC is necessary, but it is not sufficient.

We capture a real SERP snapshot for priority query clusters to understand:

- What the top results promise in their titles/snippets.
- Which SERP modules show up (PAA, video packs, forums, AI summaries).
- What page formats are winning (lists, tutorials, tools, calculators, templates).

SERPs vary by device, locale, and personalisation. Treat each snapshot as a directional sample, not an eternal truth.

Input C: Freshness checks (live sources)

Before we ship a brief, we check:

- Is the "official" source updated?
- Did the product/policy/tool change recently?
- Are there newer examples we should cite?

For fast-moving AI tooling, this step prevents the worst failure mode: publishing a confident, wrong page.

3 The weekly loop (step by step)

Step 1: Pull a GSC window and pick your battleground

We usually look at two windows:

- Last 28 days: what is moving now.
- Last 90 days: what is reliably large.

Then we pick one battleground cluster at a time (one niche, one product line, one topic tree).

Step 2: Bucket opportunities

We classify every candidate page/query into one bucket:

1. CTR wins (high impressions, good rank, low CTR).
2. Striking distance (avg position roughly 4-20 with real impressions).
3. Cannibalisation (multiple pages showing for the same intent).

4. Technical cleanup (variants, redirect/canonical gaps, soft failures).
5. Topical expansion (adjacent intents with no dedicated page yet).

Step 3: Capture a SERP snapshot for the winners

For each priority query cluster, we capture:

- Top organic results.
- Snippet patterns.
- Recency cues ("Updated", year tags).
- Content structure (TOC, FAQs, tools, downloadable assets).

We do this with Playwright (browser automation) because:

- It forces us to see the SERP the way users see it.
- It makes the capture reproducible.
- It helps us avoid building a plan on stale assumptions.

Tooling note: we do this with Playwright MCP. If you are doing it manually, an incognito window with explicit locale settings is still better than guessing.

Step 4: Run a freshness check before writing

We do a quick, explicit source pass:

- Find the most authoritative primary sources.
- Confirm dates and version numbers.
- Save URLs and citations to embed in the brief.

Tooling note: we use live search here to avoid stale claims (especially for AI tools that ship fast). If something is time-sensitive and you cannot verify it, rewrite the claim or omit it.

Step 5: Choose the smallest high-leverage action

Do not default to "write a new post".

If a page already ranks well but has low CTR, we start with:

- Title + meta description alignment.
- First-screen rewrite (deliver the promise fast).
- Better internal linking to the next click.

If we are in striking distance, we often need:

- One deeper section that matches competitor phrasing/structure.
- 2-3 supporting pages to strengthen topical authority.
- Clear canonical targeting (no duplicate intent pages).

4 Our prioritisation scorecard

We score each opportunity (0-5 each) and pick the top 3 to execute:

- Demand: impressions and long-tail breadth.
- Achievable gain: how close to top 5.
- Business intent: does it support high-intent actions.
- Topical fit: does it strengthen the cluster.
- Effort: research + production cost.

Example weighted formula:

```
priority_score =
  demand * 0.35 +
  achievable_gain * 0.25 +
  business_intent * 0.20 +
  topical_fit * 0.15 -
  effort * 0.05
```

5 Outputs we produce every cycle

Output A: a triage table

query_cluster	primary_url	signal	action	notes
"AI video workflow"	/blog/...	striking-distance	refresh	missing competitor phrasing
"voice cloning"	/landing/...	low impressions	internal-links	feed from TOF posts
"SERP analysis"	(new)	expansion	new page	template-heavy intent

Output B: a content brief with citations

Minimum fields:

- Target intent and query cluster.
- Why now (GSC + SERP evidence).

- Structure outline.
- Sources that must be cited.
- Internal links (inbound + outbound).
- Success metric and review window.

Output C: a change log

For each shipped change, we log:

- What changed.
- Why we expect it to move CTR/rank.
- When we will re-check in GSC.

6 Cannibalisation rules (non-negotiable)

We keep one canonical winner per intent.

If multiple pages are earning impressions for the same query intent:

- Consolidate.
- Redirect or de-emphasise duplicates.
- Tighten internal links so Google (and users) see one obvious answer.

7 How this fits Instavar's AI production stack

SEO is not a separate department. It is demand sensing.

Once we pick the next SEO priorities, we can use the same AI production stack to ship a full asset set:

- The page itself (the canonical long-form answer).
- Short-form scripts derived from the same brief.
- Visual examples and thumbnails that match the on-page promise.
- A measurement plan (what counts as a win).

8 Want help running this loop?

If you want us to run a full GSC + SERP + freshness cycle for your niche and turn it into a ranked roadmap (plus briefs), message us: