AI SEO In Detail

Your first analysis — the prompt, the cost gate, the verification

Episode 7 · DataForSEO + Claude Code
youtube.com/@AISEOInDetail

> ***AI writes, you verify. The assistant writes the code, runs it and summarizes. You check the code, check the numbers, and check one row by hand.***

For the workspace from episode 6 (cache, seo.db, CLAUDE.md with the cost-gate rule) and DataForSEO credentials from episode 3. Shop and keywords are made up (example-shop.test, coffee equipment). Prices as of Oct 2026 — check current pricing.

1 The task prompt

```
For these 10 coffee-equipment keywords, get US search volumes and the live Google results
page. Store everything in seo.db. Tell me who ranks, and which features appear.

burr grinder · espresso machine for home · hand grinder for espresso · gooseneck kettle ·
coffee scale · pour over coffee maker · milk frother handheld · espresso tamper · coffee
dripper · coffee grinder cleaning tablet

Plan first: check the cache, then show me the paid calls and the estimated cost, and wait for
my GO.
```

Expected plan: **check the cache** → **one search volume task** (all ten) → **ten live results pages**, depth 10, US/English → **parse** into keyword_volume + serp_snapshot, write a summary. Before you approve the script, answer from the code:

- ☐ Which **endpoints** does the script call?
- ☐ **How many times**?
- ☐ Where does the **raw response** go before anything is parsed?

> **Read the script first.** If you can't answer those three questions from the code, don't approve it.

2 The cost gate — four requirements before any paid call

WHAT
Show what will be called: the endpoint and its parameters (depth, location, language, device).

HOW MANY
Show the number of requests and tasks, after cache hits are subtracted.

COST
Show the estimated cost per call type and in total.

WAIT FOR GO
Stop. Nothing runs until you type GO — for this plan, not an earlier one.

```
Paid calls planned:   1 × search volume task (10 keywords)   ≈ $0.09
                     10 × live SERP, advanced, depth 10      ≈ $0.03
                     Total ≈ $0.12 · cache hits: 0 · Type GO to continue.
```

Twelve cents is nothing — the gate is about the habit: ten thousand keywords or a retry loop use the same code path. Every response is cached before parsing (keyword, location, depth, date in the file name); run it again and the gate says **10 cache hits, 0 paid calls, $0**.

3 Verify in five minutes

The depth-ten trap: depth counts **every element** on the page — an AI Overview counts, each "People also ask" box counts. On a feature-heavy page, depth 10 can hold only six to nine organic results, so "not in the top ten" can be an artefact of the request. And a site cited *inside* the AI Overview never shows up in organic rows.

- ☐ **Read the request parameters** in the code: depth, location, language, device. Most wrong conclusions start there.
- ☐ **Count the organic results per keyword.** If any page has fewer than ten, a "not in the top ten" claim isn't proven.
- ☐ **Open one raw response** from the cache yourself, and find one row of the summary in it.
- ☐ **Ask: what else could explain this?** A short page, a different location, or a feature that holds the site (an AI Overview citation).

> **Cheap fix:** re-run only the short pages at depth 20 (≈ $0.006 each) — through the same cost gate.

Narration is AI-generated; the method and the checks are real.