The 5-Layer Prompt Stack

The 5-Layer Prompt Stack is an enterprise-grade prompt engineering framework designed to safely steer Large Language Models (LLMs) through highly complex, capital-intensive, and safety-critical industries like mining and oil & gas.

Layer	Purpose	Mining & Oil Application	Heavy Industry Example
1. Persona	Establishes the authoritative AI mindset and cognitive bounds.	Routes the LLM's processing through niche heavy-industry logic.	Senior Drilling Engineer, HSE Compliance Officer, Exploration Geologist.
2. Context	Injects critical parameters, telemetry, and environmental data.	Pulls in real-world data like well logs, assay data, or legal terms.	"Analyzing data from Well X-202 with bottom-hole temperatures exceeding 150°C."
3. Task	Defines the exact operational or analytical directive.	Triggers high-leverage workflows instead of generic summaries.	Calculate ROP optimization, cross-examine safety logs, audit a Pugh clause.
4. Constraint	Locks down safety, compliance, and hallucination parameters.	Prevents unsafe technical assumptions and enforces industry standards.	"Only use historical sensor data. Flag data gaps. Do not calculate hypothetical yields."
5. Output Format	Shapes the reasoning process and structure of final assets.	Standardizes findings for engineering reviews and executive handoffs.	Markdown tables, API-ready JSON, safety alert bulletins, or standard code snippets.

All the work below the line is the work of **Weng Honn.**
AI tools only work well if the person using them already understands the basics.

How I prompt AI to screen oil and mining stocks

The exact 5-prompt framework I use before touching an oil or mining stock

“Stop asking Claude to find the ‘most promising’ oil company”.

That’s a fair question, and it’s the same one I asked myself the first time I tried using AI to screen a commodity sector. The answer isn’t “the AI is bad at this.”
The answer is that oil, gas, and mining are some of the worst sectors to prompt casually , and some of the best to prompt well.

The prompt most people get wrong about oil and mining stocks
Here’s why. A generic prompt gives you a generic answer: biggest market cap, recent headlines, a vague nod to “strong fundamentals.” But commodity companies don’t work like software companies.

Their trailing financials swing wildly with spot prices that were $95 a barrel eighteen months ago and $65 today. Ask the AI to screen on last year’s numbers and you’re really just asking “which company benefited most from last year’s oil price”, which tells you almost nothing about which company is actually well-run.

So I built a small stack of prompts , five, specifically , that force the AI to reason the way an actual commodities analyst would, instead of the way a headline-scanner would. I use it for oil and gas. It transfers almost unchanged to mining. Below is the exact stack, in order, with why each one exists.

The 5-Layer Prompt Stack

The core idea: a commodity company can't be judged on its own , it has to be judged against a stated price assumption, within its correct sub-segment, on the metrics that segment actually lives or dies by. Skip any one of those three things and the AI (or you) will end up comparing apples to a barrel of oil.
Here's the stack, one layer per prompt.

Layer 1 — Classify before you screen

Oil & gas isn't one industry. Upstream explorers, integrated majors, oilfield services, midstream pipelines, and refiners all have completely different value drivers, comparing a junior E&P to Exxon on the same scorecard is like comparing a corner bakery to a supermarket chain because they both sell bread.

> Prompt: "Classify the oil & gas industry into its core sub-segments (upstream E&P, integrated majors, oilfield services, midstream, refining). For each, tell me where we currently sit in the commodity price cycle and which segment offers the best risk-adjusted entry point right now. Verify current WTI and Brent spot prices before answering."

Why this matters: this single prompt does more work than people expect. It forces the model to commit to where in the cycle we are before it recommends anything — and it stops you from accidentally judging a refiner by an E&P company's playbook.

Layer 2 — Normalize to a stated price, not the trailing numbers

This is the layer that separates an amateur screen from a professional one. If you don't tell the AI what oil price to assume, it implicitly assumes "whatever it costs today, forever" — which is almost never true in a cyclical sector.

> Prompt: "Screen upstream E&P companies assuming Brent averages $70/bbl over the next 12 months. Rank your top 5 by free cash flow yield, breakeven oil price, and reserve life (years of proven reserves at current production rate) at that price assumption — not trailing twelve-month financials, which are distorted by the current spot price."

Why this matters: naming the scenario makes the whole exercise falsifiable and reusable. A reader with a more bearish oil view can re-run the exact same prompt with $55 instead of $70 and get a completely different, equally honest, answer.

Layer 3 — Screen on the metrics the sector actually runs on

Generic screening metrics (revenue growth, low debt) miss the entire point of an energy company. The metrics that matter here are industry-specific:

- Reserve replacement ratio and reserve life — how many years of production are actually left in the ground
- All-in sustaining cost / breakeven price — the oil price the company needs just to stay cash-flow neutral
- Net debt / EBITDA at the stated price, not trailing
- Capital discipline history — did management blow the last cycle's cash at the top, or hold the line?

I fold the last one into its own prompt, because it's the one that tells you the most about whether you can trust the company's next decision:

> Prompt: "For [company name], review the last full commodity cycle (2014–2020). Did management prioritize dividends/buybacks at the top of the cycle and preserve cash discipline in the downturn, or did they overspend on growth capex right before prices collapsed? Use this history to judge how trustworthy their current capital allocation plan is."

I ran a version of this on three names I was already watching, and honestly the most useful part wasn't the "winner" — it was watching how differently each company behaved the last time oil cratered.

How to practice this yourself:

- Pull up the last 10-K covering 2015 or 2020 and check whether the dividend was cut, held, or grown through the trough
- Cross-check management's own words in that year's shareholder letter against what they actually did with capex
- Flag any company that grew production and leverage simultaneously right before a downturn — that's the pattern that burns people

---

Layer 4 — Demand the bear case and a falsifier

Left alone, AI defaults to confident, tidy answers. The single best fix I've found is to explicitly demand the counter-argument and a specific number that would prove the pick wrong.

> Prompt: "Give me your single best oil & gas pick for the next 3 years. Then argue against it as a skeptical short-seller would — what's the strongest reason this thesis fails? Name one specific, checkable data point (a reserve report figure, a cost metric, a production number) that, if it comes in worse than expected in the next 2 quarters, would kill the thesis."

Why this matters: this is the layer that fights confirmation bias directly. It also gives you something concrete to actually check next quarter, instead of a vague "monitor the situation" — the thesis either survives contact with the next reserve report, or it doesn't.

Layer 5 — Verify before you trust any number

Commodity prices and reserve figures go stale within days. The model's training data has a cutoff; spot prices don't wait for it.

> Prompt: "Verify current spot WTI/Brent and this company's most recent quarterly reserve report before answering; flag if your data might be outdated."

I add a version of this line to the end of every prompt in this stack now, not just as a standalone step. It costs one sentence and it's the difference between a plausible-sounding analysis and one you can actually act on.

Putting it together

Here's the whole thing chained into a single prompt, the way I actually use it:

> "You are screening upstream E&P companies for a promising long-term investment. First classify the current point in the oil price cycle. Then screen using breakeven price, reserve life, and FCF yield assuming Brent averages

> $70/bbl. Give me your top 2 picks, the strongest bear case against your #1 pick, and the single data point that would falsify the thesis. Verify current prices and most recent reserve data before answering."

That one prompt does the work of five separate ones once you've internalized the stack — classify, normalize, screen, stress-test, verify.

Extending this to mining

The same framework transfers to gold and base metals mining almost unchanged — swap the segment list and the sector-specific metric, keep everything else:

> Prompt: "Apply the same screening framework to gold/base metals mining instead of oil: classify sub-segments (majors, mid-tier producers, royalty companies, junior explorers), screen by all-in sustaining cost (AISC) per ounce, reserve grade, and jurisdiction/political risk, assuming gold averages $[X]/oz. Give me a top pick, the bear case, and the falsifier."

What changes between oil and mining:

	Oil & Gas	Mining
Price anchor	WTI / Brent	Spot gold / base metal price
Cost metric	Breakeven oil price	All-in sustaining cost (AISC) per ounce
Reserve metric	Reserve replacement ratio, reserve life	Reserve grade, mine life
Extra risk factor	OPEC+ supply decisions	Jurisdiction / political risk of the mine site

Everything else — the classify-first step, the stated-price normalization, the bear case and falsifier, the live-data verification — is identical.

Where this still falls short

This framework doesn't pick winners for you. It structures the question well enough that a wrong answer becomes visible fast, rather than staying hidden behind a confident-sounding paragraph. I've had this stack talk me out of names I liked just as

often as it's pointed me toward new ones — that's the whole point of the bear-case layer.

It also won't save you from a genuinely bad oil price call. If Brent goes to $40 instead of your stated $70, every "breakeven at $70" company on your list needs re-screening — that's a feature, not a bug, since it's exactly the assumption the prompt made you name out loud in the first place.

If there's one line worth taking from this: stop asking AI which oil company is "promising" — ask it which one survives your own stated price scenario, and make it argue against its own answer before you believe it.

Sincerely, Weng

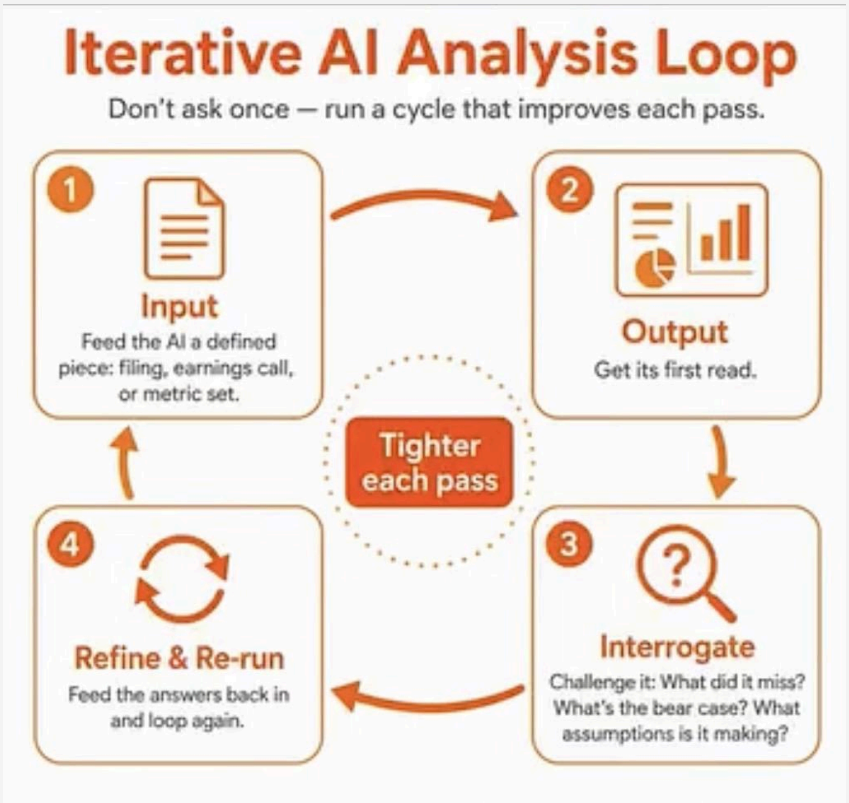

How to audit a quarterly report announcement

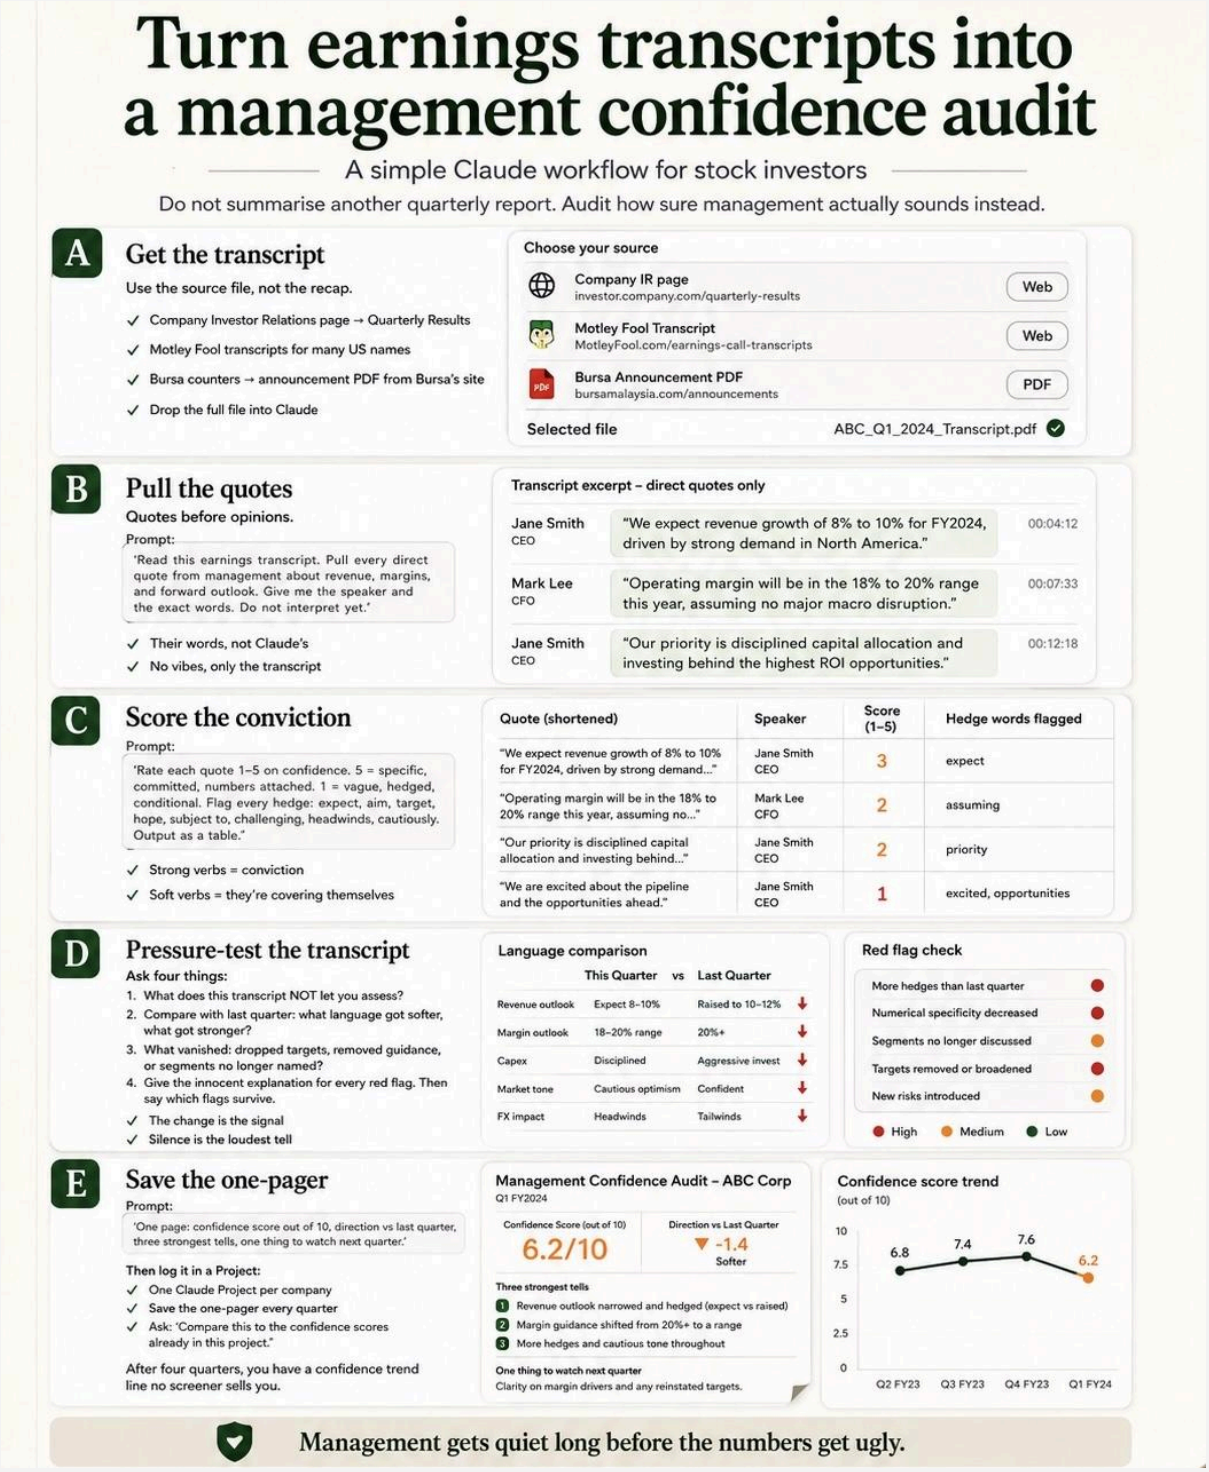

Do not summarise another quarterly report.
Audit how sure management actually sounds instead:

Get the transcript.

✅ Company's Investor Relations page → Quarterly Results (free, always there)

✅ Or Motley Fool transcripts for US names

✅ Bursa counters: the announcement PDF off Bursa's site

✅ Drop the full file into Claude. Not the media summary.

Quotes before opinions.

Prompt: "Read this earnings transcript. Pull every direct quote from management about revenue, margins, and forward outlook. Give me the speaker and the exact words. Do not interpret yet."

✅ Their words, not Claude's

✅ No vibes, only the transcript

Score the conviction.

Prompt: "Rate each quote 1-5 on confidence. 5 = specific, committed, numbers attached. 1 = vague, hedged, conditional. Flag every hedge: expect, aim, target, hope, subject to, challenging, headwinds, cautiously. Output as a table."

✅ Strong verbs = conviction

✅ Soft verbs = they're covering themselves

Make it admit the gaps.

Prompt: "Tell me what this transcript does NOT let you assess. What's missing that I'd need."

✅ Thin transcript → thin audit

✅ Garbage in = garbage out

Run the delta.

Upload last quarter's transcript too.

Prompt: "Compare these two quarters. Same topic, side by side. What language got softer, what got stronger?"

✅ The change is the signal, not the score

Hunt what vanished.

Prompt: "What did management discuss last quarter and NOT mention this quarter? List every dropped target, removed guidance, and segment that stopped being named."

✅ Silence is the loudest tell

Force it to fight back.

Prompt: "Give me the innocent explanation for every red flag. Then tell me which flags survive it."

✅ Pushback beats flattery

✅ Kill your own thesis before the market does

Demand the one-pager.

Prompt: "One page: confidence score out of 10, direction vs last quarter, three strongest tells, one thing to watch next quarter."

✅ Short enough to reread

✅ Long enough to act on

Log it in a Project.

✅ One Claude Project per company

✅ Save the one-pager every quarter

Prompt: "Compare this to the confidence scores already in this project."

That's all it takes. One honest audit.

Four quarters in, you have a confidence trend line no screener sells you.

Management gets quiet long before the numbers get ugly.