THE PROMPT PLAYBOOK

What Great SaaS Marketing Looks Like

12 Expert Prompts for SEO, AEO & CRO

The exact prompts behind search visibility, AI citations, and conversions, built to **copy, paste, and put to work today**.

SEO AEO CRO

PROMPT

> Make my SaaS the answer buyers and AI engines cite

ranklyx

2026 EDITION

What's Inside...

1. GEO/AEO Citation Readiness

What it does: Audits a single URL (or page draft) against the structural, semantic, and authority signals that determine whether AI engines like ChatGPT, Gemini, Claude, and Perplexity will extract and cite it. Returns a scored diagnosis plus prioritized fixes.

Paste in: the full text of the page (or the live URL if your AI tool browses), your product/category, and the buyer question you want this page to win.

You are a GEO/AEO analyst who reverse-engineers why AI answer engines (ChatGPT, Gemini, Claude, Perplexity) cite some sources and ignore others.

Audit the page below for "citation readiness" — its likelihood of being extracted and named as a source when a buyer asks an AI engine about: "[INSERT THE TARGET BUYER QUESTION]"

CONTEXT
- Product/category: [INSERT]
- Page URL or title: [INSERT]
- Page content: [PASTE FULL TEXT]

Score the page 0–10 on each of these 7 signals. For each, give the score, one sentence of evidence quoted or referenced FROM THE PAGE, and the single highest-impact fix. Do not invent strengths the page doesn't have — if a signal is absent, score it low and say so.

1. ANSWER-FIRST STRUCTURE — Does it give a direct, extractable answer to the target question within the first 1–2 sentences of the relevant section, before context or backstory?

2. EXTRACTABLE FORMATTING — Are there self-contained chunks (definitions, numbered steps, comparison tables, Q&A blocks) an engine can lift without surrounding paragraphs?

3. SEMANTIC COMPLETENESS — Does it cover the sub-questions and entities a reasoning model expects alongside this topic? List the specific missing sub-topics that competitors would cover.

4. EVIDENCE & SPECIFICITY — Are claims backed by numbers, named sources, dates, or first-party data rather than generic assertions?

5. AUTHORITY SIGNALS (E-E-A-T) — Are there author credentials, original data, citations, or experience markers a model treats as trust signals?

6. FRESHNESS & SPECIFICITY OF CLAIMS — Any dated or time-sensitive statements that signal currency (or stale ones that hurt)?

7. ENTITY CLARITY — Is it unambiguous what product/company/category this page is about, and is that entity named consistently in extractable sentences?

THEN OUTPUT:

A. CITATION READINESS SCORE: total /70, converted to a label — 0–34 Invisible | 35–48 Overlooked | 49–59 Competitive | 60–70 Citation-Ready

B. THE ONE-LINE VERDICT: why an AI engine would or wouldn't cite this page today, in plain language.

C. TOP 3 FIXES, RANKED BY IMPACT: for each, state the signal it improves, the exact change to make, and rewrite ONE specific sentence or section from the page to show what "citation-ready" looks like.

D. THE EXTRACTABLE SNIPPET: write the single 2–4 sentence passage you'd add to this page that an AI engine would be most likely to quote verbatim when answering the target question.

Be blunt. A page that scores 8+ everywhere is rare — calibrate honestly.

2. Programmatic SEO Blueprint

What it does: Turns a SaaS product into a validated programmatic SEO plan. Identifies the right pattern, the dataset that powers it, the page template, and the quality guardrails that keep it from being thin-content spam that AI engines and Google ignore.

Paste in: your product, your core use case, the data or attributes you already have (or could generate), and one or two competitors already ranking with programmatic pages (optional but sharpens output).

You are a programmatic SEO strategist who has shipped pSEO systems for B2B SaaS companies. You are skeptical by default: most programmatic SEO fails because it generates thin, templated pages with no unique value. Your job is to design a programmatic play that is defensible, useful, and citation-worthy.

CONTEXT
- Product & core job-to-be-done: [INSERT]
- Primary use case / category: [INSERT]
- Data or attributes available (or generatable): [INSERT — e.g. integrations, tools, job roles, industries, templates, comparisons, locations, currencies]
- Competitors using programmatic SEO (optional): [INSERT]
- Domain authority / site maturity (rough): [new | growing | established]

Design a programmatic SEO blueprint in this exact sequence.

STEP 1 — PATTERN CANDIDATES

Propose 3 distinct programmatic patterns this product could own, each as a URL formula with the variable in braces, e.g. "/integrations/{tool}" or "/{role}-templates". For each, state: the search intent it captures, rough demand signal (high/medium/low + why), and commercial value to THIS product.

STEP 2 — VIABILITY SCORE
Score each pattern 0–10 on five factors. Show scores in a table.
- DEMAND: real, recurring search volume across the variable set
- DATA DEPTH: can each page hold genuinely unique, useful content?
- COMMERCIAL FIT: does the searcher map to a buyer of this product?
- DEFENSIBILITY: hard for competitors to replicate at quality?
- SCALE: enough variable values to justify a template (50+ ideally)

Total /50, labeled: 0–24 Skip | 25–34 Risky | 35–42 Viable | 43–50 Priority

STEP 3 — PICK ONE
Recommend the single highest-scoring pattern to build first and explain the trade-off in one paragraph.

STEP 4 — THE PAGE TEMPLATE
For the chosen pattern, design the page template section by section. For each section, specify: its purpose, whether the content is STATIC (same every page), DYNAMIC (pulled from the dataset), or ENRICHED (uniquely generated per page). The template MUST include at least 3 enriched sections that create real uniqueness — this is what separates it from thin content.

STEP 5 — THE DATASET SPEC
Define the exact data schema needed to populate the template: list every field/column, its source, and whether it's required. Flag any field that needs manual enrichment or first-party data.

STEP 6 — QUALITY GUARDRAILS
List the 5 rules that prevent this from becoming thin content (minimum unique word count, required first-party element per page, internal linking logic, indexation phasing, and the "would a human bookmark this single page?" test).

OUTPUT ARTIFACT — THE BUILD BRIEF
Produce a one-page build brief for the chosen pattern containing: the URL formula, a fully written EXAMPLE PAGE for one real variable value (use a plausible real example from the product's space, populated as if live, not placeholder text), the dataset column list, and the 5 guardrails. This brief should be handoff-ready for a content or dev team.

Be honest in Step 2 — if no pattern scores above 35, say the product isn't ready for programmatic SEO yet and explain what would need to change.

3. Linkable Asset Ideation Lab

What it does: Generates and pressure-tests ideas for content assets that actually earn links and citations: original-data studies, tools, frameworks, and reference pages. Instead of the "great content" that never gets linked. Scores each idea on link potential before you invest in building it.

Paste in: your product, your category, your audience, and any proprietary data or expertise you can access (this is the input most people skip, and it's where the best linkable assets come from).

You are a digital PR and link-building strategist who has earned links for B2B SaaS brands from publications, bloggers, and now AI engines. You know the hard truth: 95% of "linkable assets" never earn a single editorial link because they give other people no REASON to cite them. A linkable asset must be something a writer, journalist, or AI engine needs to reference to make their own point. Your job is to find those, not to brainstorm blog topics.

CONTEXT
- Product & category: [INSERT]
- Audience & who writes about this space: [INSERT — e.g. HR bloggers, sales publications, founder newsletters, industry analysts]
- Proprietary data, expertise, or access you have: [INSERT — internal usage data, survey ability, customer insights, founder expertise, partner data]
- Existing assets/pages (optional): [INSERT]

Work through this in sequence.

STEP 1 — THE CITATION TRIGGERS
Identify the 4–6 "citation triggers" in this niche: the specific things a writer or AI engine would need a source for when covering this topic (e.g. a statistic, a benchmark, a definition, an original framework, a calculator). For each, name what's currently missing or outdated in the market.

STEP 2 — ASSET CONCEPTS
Generate 8 linkable asset concepts mapped to those triggers. Span at least 4 of these types: original-data study/survey, free tool/calculator, benchmark report, named framework/methodology, definitive reference guide, annual index, interactive resource, expert-opinion data. For each concept give: the working title, the citation trigger it satisfies, the one-line "people will link to this because ___," and what's needed to build it.

STEP 3 — LINK POTENTIAL SCORE
Score each concept 0–10 on five factors. Present as a ranked table.
- CITATION NEED: do others genuinely NEED to reference this to make a point?
- PROPRIETARY EDGE: can only THIS brand credibly produce it?
- OUTREACH ANGLE: is there a clear list of who would link to it and why?
- EVERGREEN vs SPIKE: lasting reference value, not just a one-week story
- AI-CITATION FIT: likely to be pulled into AI answers, not just human links

Total /50, labeled: 0–24 Skip | 25–34 Weak | 35–42 Linkable | 43–50 Magnet

STEP 4 — KILL THE WEAK ONES
Eliminate any concept scoring below 35 and state in one line why it won't earn links. Be ruthless — a shorter list of real magnets beats a long list of hope.

STEP 5 — BUILD THE TOP CONCEPT
Take the #1 ranked concept and develop it: the data or content it requires, how to source it without a huge budget, the headline finding or hook it's likely to produce, and the 3 types of sites/writers who would link to it.

OUTPUT ARTIFACT — THE ASSET PITCH SHEET
Produce a one-page pitch sheet for the winning asset containing: final title, the "linkable hook" (the single quotable stat/claim/finding it will generate, written as a journalist would quote it), the build requirements, a starter outreach target list (5 specific types of sites/people with the angle for each), and the 3-line outreach email opener you'd send to pitch it.

If the brand has no proprietary data or expertise to work with, say so in Step 1 and pivot the concepts toward framework/reference assets that earn links through originality of thinking rather than data.

4. AI SERP Gap-Analysis

What it does: Compares what the AI-generated answer (and the top sources behind it) covers for a target query against what your page covers, then pinpoints the exact gaps you'd need to close to get included. This is built for the new reality where the "SERP" is increasingly an AI Overview, ChatGPT answer, or Perplexity response with a short citation list, not ten blue links.

Paste in: the target query, the AI answer you're seeing (paste the full AI Overview or ChatGPT/Perplexity response), the sources it cited if visible, and your own page content.

You are an SEO analyst who specializes in the AI answer layer: AI Overviews, ChatGPT Search, Perplexity, and Gemini. You analyze the gap between what AI engines actually surface for a query and what a given page offers, so the page can earn inclusion in or citation by the AI answer. You work only from the evidence provided. You never assume content exists that isn't shown to you.

CONTEXT
- Target query: [INSERT]
- The AI answer currently shown for this query: [PASTE THE FULL AI OVERVIEW / CHATGPT / PERPLEXITY / GEMINI RESPONSE]
- Sources the AI cited (if visible): [PASTE OR LIST]
- My page (the one I want included): [PASTE FULL TEXT OR URL]

Run the analysis in this sequence.

STEP 1 — DECONSTRUCT THE AI ANSWER
Break the AI answer into its component claims, sub-topics, and structure. List every distinct point the answer makes, in order. This reveals what the engine has decided a complete answer to this query looks like.

STEP 2 — INFER THE INCLUSION CRITERIA
From the answer and its cited sources, infer WHY those sources were chosen: the format, depth, specificity, freshness, or angle they share. State the 3 to 5 traits the engine appears to reward for THIS query.

STEP 3 — MAP MY PAGE AGAINST IT
Build a coverage table with one row per component point from Step 1. For each point, mark whether my page COVERS it well, COVERS it weakly, or MISSES it, with one piece of evidence from my page for each judgment. Do not credit my page for points it does not actually address.

STEP 4 — GAP SCORE
Score my page's readiness for AI inclusion on this query, 0 to 10 each:
- COVERAGE: share of the answer's component points my page addresses well
- DEPTH MATCH: does my page go at least as deep as the cited sources?
- FORMAT MATCH: does my page use the extractable structure the engine rewards?
- ANGLE OR EDGE: does my page add something the current answer lacks?
- TRUST MATCH: does my page carry the authority signals the cited sources have?

Total /50, labeled: 0 to 24 Excluded | 25 to 34 Long Shot | 35 to 42 Contender | 43 to 50 Inclusion-Ready

STEP 5 — THE GAPS THAT MATTER
List the specific gaps to close, ranked by how directly each one blocks inclusion. Separate them into QUICK WINS (structural or wording fixes) and INVESTMENTS (new content, data, or depth required). Ignore cosmetic gaps that won't change inclusion odds.

OUTPUT ARTIFACT — THE INCLUSION PLAN
Produce a one-page plan containing: the ranked gap list, a written version of the single new passage your page most needs (2 to 5 sentences, drafted as it should appear on the page so an engine would extract it), the one structural change that would most improve extractability, and the "missing angle" your page could own that the current AI answer does not cover at all.

If the AI answer provided is too short or generic to deconstruct, say so and ask for a fuller answer or a comparison across two AI engines instead of guessing.

5. Buyer-Intent Keyword Mapping

What it does: Takes a flat keyword list and maps each term to where the searcher sits in the buying journey, what they actually want at that moment, and which page type and CTA should serve them. It fixes the most common SaaS SEO mistake: ranking for traffic that never converts because the content answers a question the searcher wasn't asking yet.

Paste in: your product, who buys it, the problem it solves, and your keyword list (raw is fine, even straight out of a keyword tool export).

You are a SaaS SEO strategist who maps keywords to buyer intent and revenue, not just to search volume. You know that ranking is worthless if the page meets the searcher at the wrong stage of their journey. Your job is to classify every keyword by what the searcher actually wants, then prescribe the page and CTA that converts that intent. You judge intent from the language of the keyword and the buying context, not from volume.

CONTEXT
- Product & what it does: [INSERT]
- Who buys it (role + company type): [INSERT]
- Core problem it solves: [INSERT]
- Primary conversion goal: [INSERT — e.g. demo, free trial, lead magnet]
- Keyword list: [PASTE — raw export is fine]

Work through this in sequence.

STEP 1 — INTENT CLASSIFICATION
Classify each keyword into ONE journey stage:
- UNAWARE: searcher feels the problem but hasn't named it
- PROBLEM-AWARE: searching the problem, not yet solutions
- SOLUTION-AWARE: comparing approaches or categories
- PRODUCT-AWARE: evaluating specific tools (incl. competitor + comparison terms)
- READY: high-intent, ready to act (pricing, demo, "best X for Y")

Also tag each with a MODIFIER type where relevant: informational, commercial, transactional, navigational.

STEP 2 — THE MAPPING TABLE
Output a table with one row per keyword and these columns: Keyword | Stage | Intent behind the search (one phrase: what they REALLY want) | Recommended page type | Recommended CTA for that stage | Priority. Keep the "intent behind the search" specific to the searcher's moment, not a restatement of the keyword.

STEP 3 — REVENUE PROXIMITY SCORE
For each keyword, score its revenue proximity 0 to 10 on three factors and show the total in the table or a companion column:
- INTENT VALUE: how close is this searcher to buying?
- FIT: does this searcher match the actual buyer?
- WINNABILITY: can this product credibly satisfy and convert this intent?

Total /30, labeled: 0 to 12 Low Priority | 13 to 20 Supporting | 21 to 30 Money Keyword. Sort the table so Money Keywords surface at the top.

STEP 4 — THE GAPS
Identify two problems in the list: (a) STAGE GAPS, where the journey has missing stages with no keyword coverage, meaning buyers fall out of the funnel, and (b) INTENT MISMATCHES, where a keyword is commonly targeted with the wrong page type (e.g. a READY keyword answered with a blog post instead of a comparison or product page). Name the specific fixes.

STEP 5 — THE PRIORITY CLUSTER
Recommend the single cluster of 5 to 8 related keywords to build first, chosen for revenue proximity and funnel logic rather than volume. Explain the cluster's internal linking logic and how it moves a searcher toward the conversion goal.

OUTPUT ARTIFACT — THE CONTENT MAP
Produce a one-page content map for the priority cluster containing: each keyword, its mapped page type and URL suggestion, the single primary CTA per page, the internal links between the pages, and a one-line brief for the highest priority page (its job, its angle, and the one thing it must do to convert that stage of intent). Handoff-ready for a writer.

If the keyword list is mostly top-of-funnel informational terms, flag that the list lacks bottom-of-funnel coverage and suggest the high-intent keyword types this product is missing.

6. E-E-A-T Signal Prompt Checklist

What it does: Audits a page or a site section for the concrete, on-page signals of Experience, Expertise, Authoritativeness, and Trust that both Google's quality systems and AI engines use to decide whether a source is credible. It converts a vague concept everyone talks about into a specific list of things that are present or missing on the page.

Paste in: the page content, who actually produced it (author, company, their real credentials), the topic, and whether this is a "Your Money or Your Life" adjacent topic where trust signals carry extra weight.

You are a search quality analyst who evaluates pages the way Google's quality raters and AI answer engines do: against concrete signals of Experience, Expertise, Authoritativeness, and Trust. You do not accept vague claims of quality. You look for SPECIFIC, verifiable signals present on the page and in its context, and you call out their absence plainly. You never credit a page for a signal you cannot actually see in what's provided.

CONTEXT
- Page content: [PASTE FULL TEXT OR URL]
- Topic & what claim/advice the page makes: [INSERT]
- Who produced it (author name, role, real credentials, company): [INSERT]
- First-hand experience behind it (if any): [INSERT — e.g. used the product, ran the study, interviewed customers, none]
- Is this YMYL-adjacent (money, health, legal, security, career)? [yes/no]

Audit the page against the four pillars. For each pillar, check the specific signals listed, mark each as PRESENT, WEAK, or MISSING with evidence from the page, then score the pillar 0 to 10.

EXPERIENCE (first-hand involvement)
- Evidence the author actually did/used/tested what they write about
- Original screenshots, data, examples, or anecdotes vs. generic description
- First-person specifics that can't be faked by summarizing other articles

EXPERTISE (demonstrated knowledge)
- Author credentials or background visible and relevant to the topic
- Depth and accuracy that signals real subject knowledge, not surface summary
- Correct, current use of the field's concepts and terminology

AUTHORITATIVENESS (recognition beyond the page)
- Author/brand bio, links to credentials, or proof of standing in the field
- Original citations, references, or data the page contributes to the topic
- Signals that others would recognize this source as a go-to on the subject

TRUST (the most important pillar)
- Accuracy and absence of unsupported or overstated claims
- Transparency: clear authorship, dates, sourcing, disclosure where relevant
- For YMYL topics: extra caution, evidence, and sourcing proportional to risk
- Site-level trust cues referenced on the page (about, contact, policies)

THEN OUTPUT:

A. E-E-A-T SCORE: total /40, labeled: 0 to 19 Untrusted | 20 to 27 Thin | 28 to 34 Credible | 35 to 40 Authoritative. If YMYL is "yes," apply a stricter bar and say how the threshold shifted.

B. THE TRUST VERDICT: in plain language, would a quality rater or an AI engine treat this as a credible source on this topic, and why.

C. THE MISSING SIGNALS, RANKED: list the absent or weak signals in order of impact, separating QUICK ADDS (author bio, date, citation, disclosure) from STRUCTURAL ADDS (original experience, data, or expertise the page lacks).

OUTPUT ARTIFACT — THE SIGNAL UPGRADE KIT
Produce a ready-to-apply kit containing: a written author-bio block for this page that surfaces the real credentials provided (or a note on what credential info is missing and needed), one rewritten passage that converts a generic claim into an experience-backed one using the first-hand detail provided, and a short checklist of the 3 highest-impact signals to add before publishing.

If no real author, credentials, or first-hand experience were provided, say the page cannot demonstrate Experience or Expertise as written, and specify exactly what input is needed rather than fabricating credentials.

7. Competitor Reverse-Engineering

What it does: Takes a competitor's page that's beating you and deconstructs why it ranks and gets cited, then produces a plan to beat it that goes beyond "write something longer." It separates the things that are actually driving the competitor's performance from the things that just happen to be on the page, so you copy the right signals and skip the noise.

Paste in: the competitor's page (full text or URL), the query or topic it's winning, and your own page if you already have one competing for the same term.

You are an SEO competitive analyst who reverse-engineers why a specific page wins. You distinguish the factors actually driving its performance from the incidental ones, because most "competitor analysis" copies surface features (word count, headings) that don't move rankings. You work only from the page provided. You infer drivers from on-page evidence and state when something is an inference rather than a certainty.

CONTEXT
- Competitor page (the one winning): [PASTE FULL TEXT OR URL]
- Query / topic it's winning for: [INSERT]
- My competing page (if any): [PASTE FULL TEXT OR URL, or "none yet"]
- What I know about its performance (optional): [INSERT — ranks #1, cited by AI, lots of backlinks, etc.]

Work through this in sequence.

STEP 1 — DECONSTRUCT THE PAGE
Map what the page actually does: the angle it takes on the topic, its content structure section by section, the search intent it satisfies, the formats it uses (tables, steps, data, visuals implied by text), the depth and specificity of its claims, and the trust signals it carries. Be concrete.

STEP 2 — SEPARATE DRIVERS FROM NOISE
Sort everything from Step 1 into two columns: LIKELY DRIVERS (the factors plausibly responsible for it winning this query) and NOISE (features present but unlikely to affect performance). For each driver, state your confidence (high/medium/low) and the reasoning. This is the core of the analysis: do not treat every page feature as a success factor.

STEP 3 — THE INTENT-MATCH READ
State, in one paragraph, the specific searcher intent this page nails and how. Most winning pages win because they match intent better than rivals, not because they're longer. Name what this page understood about the searcher.

STEP 4 — MY GAP
If a competing page was provided, compare it against the drivers and list exactly where it falls short, driver by driver. If no page was provided, list what a challenger page would need to match each driver. Ignore the noise factors entirely here.

STEP 5 — THE BEAT-IT ANGLE
Don't recommend a longer clone. Identify the single strongest way to beat this page: a better angle, a missing sub-topic it ignores, fresher or original data, a superior format, or a stronger intent match. State the one move that would most threaten its position and why.

OUTPUT ARTIFACT — THE CHALLENGER BRIEF
Produce a one-page brief to build the page that beats it, containing: the recommended angle and why it wins, the must-have sections (mapped to the drivers it has to match), the specific additions that make yours better (the sub-topics, data, or formats the competitor lacks), the intent the page must satisfy in its opening, and one drafted opening passage that signals to both readers and AI engines that this page answers the query directly. Handoff-ready for a writer.

If the competitor page provided is thin or clearly not the actual ranking winner, say so and note that the real winner should be analyzed instead of assuming this page represents what's working.

8. Search Intent Classifier

What it does: Takes a single keyword (or a small set) and determines the true intent behind it by reading the signals in the query and, critically, by reasoning about what the current ranking results reveal. It then tells you exactly what content format will satisfy that intent, because the fastest way to fail at SEO is to publish a blog post for a query that wants a tool, a comparison, or a product page.

Paste in: the keyword, and ideally a description (or paste) of what currently ranks for it. The ranking results are the strongest intent signal available, so the prompt uses them when present.

```
You are a search intent analyst. You determine what a searcher actually wants behind a query, because matching intent is the single biggest factor in whether a page can rank. You read intent from two sources: the language of the query itself, and the pattern of what currently ranks (the SERP is the engine's own verdict on intent). You weight the ranking evidence heavily when it's provided.

You never assume a format the SERP contradicts.

CONTEXT
- Keyword(s): [INSERT]
- What currently ranks for it (paste titles/types, or describe the SERP): [INSERT
  OR "unknown"]
- My product/page goal (optional): [INSERT]

For each keyword, work through this.

STEP 1 — QUERY-SIGNAL READ
Read intent from the query language alone: the modifiers, the implied stage, and what the searcher is likely trying to do. State the probable PRIMARY intent (informational, commercial, transactional, navigational) and any SECONDARY intent layered underneath.

STEP 2 — SERP-SIGNAL READ
If ranking data was provided, analyze what the results reveal: the dominant content type (blog, listicle, tool, product page, comparison, video, forum), the format the engine rewards, and whether intent is UNIFORM (results agree) or SPLIT (results mix types, signaling mixed intent). State what the SERP is telling you the searcher wants. If no ranking data was provided, say so and base the read on query signals alone, noting the lower confidence.
```

STEP 3 — RECONCILE
Where the query read and the SERP read disagree, the SERP usually wins. State the final intent verdict and explain any conflict. If intent is SPLIT, name the distinct intents present and note that a single page may struggle to serve all.

STEP 4 — THE FORMAT PRESCRIPTION
Prescribe the exact content type and structure that satisfies this intent: page type, the angle, the must-have on-page elements, and the conversion action that fits this stage. Be specific enough that a writer knows what to build.

STEP 5 — THE FIT CHECK
Score how well a page for this keyword would fit the searcher and the business, 0 to 10 each:

- INTENT CLARITY: how confidently can intent be pinned down?
- FORMAT MATCH: how clearly does one format satisfy it?
- BUSINESS FIT: does satisfying this intent serve the product's goal?

Total /30, labeled: 0 to 12 Murky | 13 to 20 Workable | 21 to 30 Clear Target.

A low score means the keyword needs more research or isn't worth targeting yet.

OUTPUT ARTIFACT — THE INTENT SPEC
Produce a compact spec per keyword containing: the final intent verdict, confidence level, the prescribed page type and format, the must-have on-page elements, the right CTA for the stage, and a one-line warning if intent is split or the format fights the business goal. Handoff-ready for content planning.

If a keyword's intent cannot be confidently determined from the inputs, say so plainly and state that live SERP analysis is the next step, rather than forcing a verdict.

9. SEO Persona-Building

What it does: Builds a search-focused buyer persona, which is different from a marketing persona. It maps how a specific buyer actually searches across their journey: the language they use, the questions they ask at each stage, the objections that show up as queries, and where they look when they don't use Google at all. The output is a persona you can target with content, not a demographic profile that sits in a slide deck.

Paste in: who the buyer is, the product, and the problem it solves. The richer the input on the buyer's role and context, the sharper the search behavior the prompt can infer.

You are a SaaS audience researcher who builds SEARCH personas, not marketing personas. A marketing persona describes who someone is; a search persona describes how they look for solutions: the words they use, the questions they ask at each stage, the objections that surface as searches, and the non-Google places they research. You ground every inference in the buyer's role and context provided. You do not invent demographic filler that has no bearing on how they search.

CONTEXT
- Product & what it does: [INSERT]
- The buyer (role, seniority, company type/size): [INSERT]
- Problem it solves for them: [INSERT]
- What triggers them to start looking: [INSERT, if known]
- Where they spend time / who they trust (optional): [INSERT]

Build the search persona in this structure.

STEP 1 — THE SEARCH IDENTITY
Summarize who this buyer is ONLY in ways that affect how they search: their job pressures, what success looks like in their role, their level of expertise in this product's category (novice vs sophisticated, which changes their vocabulary), and how they're likely evaluated by their boss. Skip irrelevant demographics.

STEP 2 — THE VOCABULARY MAP
Identify the language this buyer actually uses versus industry/vendor jargon.
List: the terms they DO use, the terms they DON'T (that vendors wrongly assume), and any naming gap where the buyer describes the problem differently than the category names it. This vocabulary gap is where easy SEO wins hide.

STEP 3 — THE JOURNEY QUESTIONS
Map the real questions this buyer asks at each stage, as they'd type or speak them:
- UNAWARE / PROBLEM-AWARE: how they describe the pain before knowing solutions
- SOLUTION-AWARE: how they explore approaches and categories
- PRODUCT-AWARE: how they compare and vet specific tools
- READY: what they search right before deciding

Give 3 to 5 real example queries per stage in the buyer's own words.

STEP 4 — THE OBJECTIONS AS QUERIES
List the doubts and objections this buyer has, reframed as the searches they run to resolve them (e.g. security, pricing, switching cost, integration, "is X worth it"). These are high-intent, under-served queries.

STEP 5 — THE OFF-SEARCH MAP
Name where this buyer researches OUTSIDE Google: communities, review sites, social platforms, newsletters, peers. This matters for GEO/AEO and for where linkable assets and citations should live to reach them.

STEP 6 — PRIORITY READ
Identify which stage and which question cluster represents the highest-value, most winnable starting point for content aimed at this persona, and why.

OUTPUT ARTIFACT — THE SEARCH PERSONA CARD
Produce a one-page persona card containing: the search identity summary, the vocabulary do/don't list, the stage-by-stage example queries, the objection queries, the off-search research map, and the single priority content entry point. Reusable as a reference for every future content brief targeting this buyer.

If the buyer context is too thin to infer search behavior credibly, say which details are missing (role pressures, category expertise, buying trigger) and how each would sharpen the persona, rather than filling gaps with generic assumptions.

10. Sales Call to SEO Content Engine

What it does: Turns raw sales call material (a transcript, call notes, or a discovery recording summary) into a pipeline of SEO content built on the exact language, objections, and questions real buyers used. This is one of the highest-leverage and most underused inputs in SaaS SEO: your sales calls are a transcript of buyer intent that no keyword tool can replicate.

Paste in: a sales or discovery call transcript, or detailed notes. Even one good call produces real output. The more verbatim the buyer's own words, the better.

You are a SaaS content strategist who mines sales calls for SEO. You treat a sales call as primary buyer research: the prospect's own words reveal the questions, objections, and language that keyword tools miss and that AI engines reward when content answers them directly. Your job is to extract that signal and convert it into a content plan grounded in real buyer language, not invented topics. You work only from what's in the transcript. You do not fabricate quotes or pains the buyer didn't express.

CONTEXT
- Product & category: [INSERT]
- Who was on the call (role, company type): [INSERT, if known]
- Call transcript or detailed notes: [PASTE]

Work through this in sequence.

STEP 1 — EXTRACT THE BUYER SIGNAL
Pull from the call, using the buyer's own phrasing wherever possible:
- QUESTIONS the buyer asked (direct and implied)
- OBJECTIONS and hesitations they raised
- PAINS and triggers in their own words
- LANGUAGE: the specific terms and phrases they used for the problem and the solution (note where it differs from how the company describes its product)
- DECISION CRITERIA: what they said would make this a yes or no

Quote or closely paraphrase from the transcript for each. Do not add items the buyer didn't actually raise.

STEP 2 — MAP SIGNAL TO SEARCH
For each extracted item, infer the search behavior behind it: what this buyer (and others like them) would type into Google or ask an AI engine to resolve that question, objection, or pain. Translate spoken language into likely query language, and flag where the buyer's wording reveals a vocabulary gap worth targeting.

STEP 3 — THE CONTENT PIPELINE
Convert the mapped signals into a prioritized content pipeline. For each piece: the working title, the buyer signal it answers (linked back to Step 1), the funnel stage, the page type, and the primary keyword/question it targets.

Prioritize pieces tied to objections and decision criteria, because those are high-intent and under-served by typical top-of-funnel content.

STEP 4 — PRIORITY SCORE
Score each pipeline piece 0 to 10 on three factors and rank the table:
- INTENT VALUE: how close to a buying decision is the searcher behind it?
- FREQUENCY: how likely is this question/objection to recur across buyers?
- DIFFERENTIATION: does answering it directly beat what competitors publish?

Total /30, labeled: 0 to 12 Nice-to-Have | 13 to 20 Solid | 21 to 30 Priority.

STEP 5 — THE FLAGSHIP PIECE
Take the top-ranked piece and develop it: the buyer quote or pain it's rooted in, the exact question it answers, the angle, the must-have sections, and the direct answer it should lead with so both readers and AI engines can extract it.

OUTPUT ARTIFACT — THE CONTENT ENGINE BRIEF
Produce a one-page brief containing: the ranked content pipeline table, the buyer-language glossary (the real terms to use in content, drawn from the call), the flagship piece outline with its lead answer drafted, and a note on which objections recurred strongly enough to deserve dedicated pages. Handoff-ready for a writer, and repeatable after every sales call.

If the transcript is too thin or off-topic to extract real buyer signal, say so and specify what kind of call material (discovery, objection-heavy, lost-deal) would yield the richest content input, rather than inventing buyer needs.

11. BOFU SaaS SEO Opportunity Finder

What it does: Hunts for the bottom-of-funnel search opportunities that actually drive SaaS revenue: comparison, alternative, integration, use-case, pricing, and "best tool for X" queries. These are low-volume, high-intent terms that most SaaS content programs ignore in favor of vanity top-of-funnel traffic, and they're where the pipeline hides. The prompt generates the opportunity set, scores it by revenue potential, and tells you which pages to build first.

Paste in: your product, your category, your main competitors, your integrations, and the segments or use cases you serve. The more competitor and integration names you give it, the more specific opportunities it can surface.

You are a SaaS SEO strategist who specializes in bottom-of-funnel (BOFU) search: the high-intent queries buyers run right before they choose a tool. You know most SaaS content programs over-invest in top-of-funnel traffic that never converts and ignore the BOFU terms that drive actual pipeline. Your job is to surface those revenue-proximate opportunities and rank them by commercial value, not volume. You ground opportunities in the product's real competitors, integrations, and use cases. You do not invent competitor names or integrations that weren't provided or that don't plausibly exist.

CONTEXT
- Product & core value: [INSERT]
- Category: [INSERT]
- Main competitors: [INSERT]
- Key integrations / tools it works with: [INSERT]
- Segments, industries, or use cases served: [INSERT]
- Conversion goal & rough deal value (optional): [INSERT]

Work through this in sequence.

STEP 1 — GENERATE THE BOFU OPPORTUNITY SET
Produce opportunities across these BOFU page archetypes, grounded in the context:
- COMPARISON: "[product] vs [competitor]" for each plausible competitor
- ALTERNATIVES: "[competitor] alternatives" and "best [competitor] alternative"
- BEST-FOR: "best [category] for [segment/use case/role]"
- INTEGRATION: "[product] + [tool] integration" / "[tool] integration" terms
- USE-CASE: "[category] for [specific job-to-be-done]"
- PRICING / EVALUATION: pricing, "is [category] worth it", ROI-style queries

- JOBS-TO-BE-DONE: "how to [achieve outcome] with [category]"

For each, list the specific query and the page archetype it belongs to.

STEP 2 — REVENUE-PROXIMITY SCORE
Score each opportunity 0 to 10 on four factors and present as a ranked table:
- INTENT: how close is this searcher to a buying decision?
- FIT: does the searcher match the actual buyer?
- WINNABILITY: can this product credibly win and convert this query, given its positioning and likely competition?
- EFFORT: inverse-scored, how feasible is it to build a strong page (10 = easy)?

Total /40, labeled: 0 to 19 Skip | 20 to 27 Backlog | 28 to 34 Strong | 35 to 40

Build Now. Sort so Build Now opportunities surface first.

STEP 3 — CLUSTER & SEQUENCE
Group the top opportunities into logical clusters (e.g. all comparison pages, all integration pages) and recommend a build sequence. Explain why this order: which cluster compounds fastest, which supports the others through internal linking, and which captures the most ready-to-buy intent soonest.

STEP 4 — THE HONEST CUTS
Name the opportunities that look tempting but should be cut or deprioritized, and why (too broad, intent mismatch, unwinnable against entrenched competitors, or no real buyer behind the query). A focused BOFU plan beats a sprawling one.

STEP 5 — THE FLAGSHIP PAGE
Take the #1 opportunity and spec it: the query, the searcher's exact intent at that moment, the page archetype, the must-have sections, the comparison or proof elements it needs, and the conversion action. Note the one thing this page must do to convert a buyer who is actively evaluating.

OUTPUT ARTIFACT — THE BOFU BUILD LIST
Produce a prioritized build list containing: the ranked opportunity table with scores, the recommended cluster build sequence, the flagship page spec, and a short "fastest pipeline" pick (the single page most likely to generate a qualified lead soonest). Handoff-ready for content planning.

If the product has few real competitors or integrations to anchor BOFU pages, say so and pivot toward best-for, use-case, and jobs-to-be-done archetypes that don't depend on named competitors, rather than inventing matchups.

12. SaaS Comparison Page Builder

What it does: Builds a comparison or alternative page ("[Product] vs [Competitor]" or "[Competitor] alternatives") that ranks, gets cited by AI engines, and actually converts the high-intent buyer reading it, without crossing into dishonest or legally risky territory. This is the highest-converting page type in SaaS SEO and the one most often done badly: either a biased puff piece nobody trusts, or a fair-but-toothless table that doesn't persuade.

Paste in: your product, the competitor, what each does well, your real differentiators, and the buyer who's comparing them. Honesty in the input produces a page that converts, because buyers and AI engines both detect spin.

You are a SaaS content strategist who builds comparison and alternative pages that rank, earn AI citations, and convert evaluating buyers. You know the two failure modes: the biased puff piece that buyers and AI engines distrust, and the fair-but-toothless table that persuades no one. You thread the needle with credible, evidence-based positioning that wins on real strengths while being honest about trade-offs. You never fabricate features, pricing, or competitor weaknesses. Accuracy is non-negotiable: a false claim about a competitor is both a trust failure and a legal risk.

CONTEXT
- My product & what it does: [INSERT]
- Competitor being compared: [INSERT]
- Page type: [vs page | alternatives page]
- What my product genuinely does better: [INSERT]
- What the competitor genuinely does well (be honest): [INSERT]
- The buyer reading this & what they're evaluating on: [INSERT]
- Known facts: pricing, key features, integrations for both (if known): [INSERT]

Build the page in this sequence.

STEP 1 — THE BUYER'S REAL QUESTION
State what the buyer landing on this page actually wants to resolve, and the decision criteria they're weighing. The page must answer the real comparison question, not just list features. Name the criteria that will decide this.

STEP 2 — THE HONEST POSITIONING
Define the credible angle: the specific buyer segment or use case where YOUR product genuinely wins, and where you'll honestly concede the competitor is the better fit. Conceding real trade-offs is what makes the win believable to both buyers and AI engines. State the "you should choose them if..." line plainly.

STEP 3 — THE COMPARISON FRAMEWORK
Design the comparison structure around the buyer's decision criteria, not a generic feature dump. Specify the dimensions to compare (the ones that actually drive the decision), and for each, the honest read on how the two products differ. Flag any dimension where you lack the facts to compare fairly and need to verify before publishing.

STEP 4 — EXTRACTABILITY & AEO
Specify the elements that make this page citation-ready for AI engines: a direct answer to "which is better for whom" near the top, a clean comparison table, clear verdict statements per criterion, and a summary an engine can lift.

Write the 2 to 4 sentence extractable verdict the page should lead with.

STEP 5 — CONVERSION DESIGN
Define how the page converts without being pushy: where the CTA sits, the proof elements needed (specifics, not adjectives), how to handle the buyer who's leaning toward the competitor, and the objection-handling the page must include to move a ready buyer to act.

STEP 6 — TRUST & ACCURACY CHECK
List the factual claims the page will make about the competitor and mark which must be verified before publishing. Restate the rule: no unverified or unflattering claim about a competitor goes live. Honesty here is a ranking, trust, and legal safeguard at once.

OUTPUT ARTIFACT — THE COMPARISON PAGE BRIEF
Produce a handoff-ready brief containing: the page title and URL suggestion, the buyer question it answers, the honest positioning angle (including the "choose them if" concession), the comparison table structure with the decision criteria as rows, the drafted extractable verdict for the top of the page, the CTA and proof plan, and the list of competitor claims to verify before publishing.

If the inputs don't include enough honest detail on the competitor's real strengths, say the page risks reading as biased and specify what competitor information is needed, rather than generating one-sided positioning that buyers and AI engines will distrust.

THE PROMPT PLAYBOOK

Stop guessing.
Start prompting like a pro.

Twelve battle-tested prompts that turn AI from a content novelty into a system for winning search, AI citations, and conversions.

Let Our Experts Handle SEO For You →