

TAU FIELD GUIDE | SEPTEMBER 2026

THE ORGANIC SEARCH PLAYBOOK

From one search intent to a page worth clicking — or a refusal that says what would have to be true first.

10 bounded skills | BRIEF → READ → DECIDE → DRAFT → PLACE → CHECK → REVIEW



Search ranking is not a published algorithm. This playbook separates what the engines document from what the field merely believes, and refuses to publish a page that cannot earn its URL.

01 | HOW THE PACK RUNS

Order beats effort. Check before you write.

Most search work fails in one of three places: it starts from a keyword instead of a job; it writes before it checks, so a header at the CDN edge quietly wastes a week of drafting; or it cannot tell a documented fact from a widely repeated one. The order below closes all three.

#	SKILL	THE QUESTION IT ANSWERS	ENDS RUN?
1	Search brief	Do we know enough to start?	yes
2	Results-page read	What does the engine currently reward here?	yes
3	Page or no page	Should this page exist at all?	yes
4	Page contract	Does the draft keep one promise?	yes
5	Placement	How is the page reached?	no
6	Crawl + index	Can it be fetched and indexed?	yes
7	Rendering	Is the content in what the crawler gets?	yes
8	Structured data	Does a rich result type actually apply?	no
9	Page experience	Is speed a floor we clear?	no
10	Evidence + review	What may we claim, and when do we look again?	yes

THE RULE ABOVE THE REST

The page has to be worth reading if search engines did not exist. Everything in this playbook is about making such a page findable. Nothing in it is about making an unreadable page rank. Google's own framing of the same idea: content made primarily to attract search visits is not what its systems seek to reward.

SKILLS 6, 7, 9 AND 10 RUN BEFORE THE DRAFTING

- Each of them can silently make the writing worthless, and finding that out after the draft is the expensive order.
- Five of the ten skills can end the run. Refusal is a first-class output, and the cheapest place to stop is before the writing.
- A complete run hands back five artefacts: the brief, the decision, the draft, the placement plan and the check log.

PRIMARY SOURCES

[Search Essentials](#)
[Helpful content](#)
[SEO starter guide](#)

02 | THE BRIEF

Start from a job, not from a string.

The skill does not start from a keyword. It starts from a filled brief, and it stops and asks when a field is missing rather than inventing it. Six fields, all required.

TARGET INTENT

One sentence naming the job the searcher is trying to finish. “Compare two ways of doing something and pick one” is an intent. A two-word phrase is not.

SEED QUERIES

The handful of phrasings real people use, including the clumsy ones. The range is what tells you whether this is one page or three.

SITE CONTEXT

What the site already publishes on this intent, what it is credibly expert in, and which sections crawlers reach easily today.

FIRST-PARTY

Own data, own method, own product behaviour, own worked examples. A page with none of this has nothing another page cannot copy in an afternoon.

CONSTRAINTS

Claims the page may not make, names it may not use, and the review path for anything regulated or legally sensitive.

DEFINITION OF DONE

The job the page does once it ranks: a signup, a support deflection, a citation. A page with no defined job never gets improved.

COPY-READY FIRST RESPONSE

Produce the six-field brief from what I gave you. For every field you can fill, fill it and mark it CONFIRMED. For every field you cannot, do NOT invent a value: mark it MISSING and write the single question you would ask to fill it. Write the target intent as a sentence beginning “Someone who searches this is trying to”. End with a go / no-go line: GO only if all six fields are CONFIRMED.

PRIMARY SOURCES

[Helpful content](#)
[Spam policies](#)
[Search Console](#)

03 | READ THE RESULTS PAGE

The results page is the brief. It expires.

Before a word is written, look at what the engine currently returns for the seed queries. It is the closest thing to a public statement about what the engine presently treats as a good answer — and it is evidence with a date on it, so it gets re-read rather than remembered.

01

THE INTENT CLASS

Informational, navigational, commercial or transactional — visible in the kind of pages returned, not in the words of the query.

02

THE PAGE TYPE WINNING NOW

A guide, a product page, documentation, a forum thread, a video. Then: how far is that from what you planned to write?

03

THE RESULT FEATURES

A query answered by a snippet, a definition panel, an AI summary or a local pack may never pass a click to any page at all.

04

THE RECENCY SKEW

Winners that skew recent tell you the answer has a shelf life, and the page needs a named owner rather than a publish date.

WHEN THIS STEP ENDS THE TASK

If the winning format is one the site cannot credibly produce — an independent review corpus, a marketplace listing, a dataset the organisation does not have — the correct output of this step is not a page. It is a note explaining what would have to be true first.

WHAT THIS STEP MUST NOT DO

- Turn an observation into a mechanism. “The top three all have a table of contents” is an observation; “a table of contents helps you rank” is folklore.
- Record a result without its date, locale, language and device. Results differ across all four, and they decay.

PRIMARY SOURCES

Ranking systems

Snippets

AI features

Stop the site competing with itself.

Two pages competing for one intent do not split the traffic evenly and do not add up. The engine picks one — often not the one you meant — and the loser keeps drawing internal links away from the winner. Google states plainly that duplicate content is not a penalty; it is inefficiency you pay for twice.

IMPROVE

The site already appears for any seed query, even far down. An indexed page with history is a stronger starting position than a new URL.

NEW

The intent is genuinely distinct — and you can name the queries the old page keeps. If you cannot name them, you are building your own competitor.

MERGE

Two pages already answer the same question. Fold the weaker into the stronger and redirect it, rather than leaving both up.

DECLINE

The site has no standing to answer. A legitimate and frequently correct output. Hand back the reasoning, not a draft.

THE CONSOLIDATION SIGNALS GOOGLE DOCUMENTS, STRONGEST FIRST

SIGNAL	GOOGLE'S OWN CHARACTERISATION
Redirect	A strong signal that the redirect target should become canonical
rel="canonical" link	A strong signal that the specified URL should become canonical — a signal, not a directive
rel="canonical" header	The same, for non-HTML files such as PDFs
Sitemap inclusion	A weak signal
HTTPS, hreflang clusters	Additional automatic preferences

None of them is required, and the engine may still choose a different URL than the one you declared. Design the URL space so the question does not arise.

PRIMARY SOURCES

- Consolidate duplicate URLs
- Ranking systems

05 | THE PAGE CONTRACT

Seven clauses. A break blocks publication.

A page that satisfies this contract is not guaranteed to rank, because nobody outside the engine can guarantee that. A page that breaks it should not be published.

01

ONE PAGE, ONE INTENT

The page finishes one job. A second job gets a second page and a link between them.

02

THE ANSWER COMES FIRST

The direct answer sits above the background and the caveats. The page stops depending on the reader's patience.

03

THE TITLE STATES THE PROMISE

Title element and first heading make the same promise, in the searcher's vocabulary, not internal product language.

04

HEADINGS ARE THE NEXT QUESTIONS

Subheadings are the sub-questions a reader actually has next, in the order they arrive.

05

PASSAGES STAND ALONE

Any section survives being copied into an empty document: subjects named, terms defined where used, no callbacks.

06

CLAIMS ARE ATTRIBUTABLE

Every number, date and quotation has a reachable source, and first-hand experience says so plainly.

07

NOTHING LOAD-BEARING BEHIND A SCRIPT

What the page is about belongs in the HTML the server sends. Content that only appears after a client-side fetch may be indexed late, partially, or not at all. Skill 7 tests this properly; at drafting time the rule is simply to raise it now.

Google states that heading order does not matter from a search perspective, that there is no magic word count, and that the keywords meta tag is not used. Spend the review cycle on clauses 2, 5 and 6 instead — those are the ones readers feel.

PRIMARY SOURCES

[Title links](#)[Helpful content](#)[SEO starter guide](#)

Two jobs, routinely confused.

Neither is used verbatim. Title link generation is automated and considers both the content of a page and references to it elsewhere on the web; snippets are created automatically from page content, and the meta description is used when it gives users a more accurate description.

	<title> ELEMENT	META DESCRIPTION
What it feeds	The title link, as one input among several	Sometimes the snippet
Written for	The searcher's vocabulary; states the promise	A short, relevant, accurate summary of this page
Failure mode	Boilerplate, stuffed, or promising a different page	Identical site-wide; written for a robot
Length	Descriptive and concise; no documented limit	No documented limit; write for accuracy, not a pixel target

THE SOURCES GOOGLE DOCUMENTS FOR A TITLE LINK

- The <title> element, the main visual title on the page, and heading elements such as <h1>.
- og:title, prominently styled text, anchor text in links pointing at the page, and WebSite structured data.
- Best practice, verbatim in spirit: every page has a title; descriptive and concise; avoid keyword stuffing; avoid repeated or boilerplate titles; brand concisely; make it clear which text is the main title.

SNIPPET CONTROLS WORTH KNOWING

nosnippet suppresses the snippet entirely. max-snippet:[number] caps its length. The data-nosnippet attribute excludes one part of the page. All three also affect what Google's AI experiences can show from your pages, because those experiences use the same Search controls — there is no AI-only opt-out inside Search. Use them deliberately; suppressing your own snippet is rarely what you want.

PRIMARY SOURCES

- Title links
- Snippets
- AI features

A page nobody links to is reached late.

Google finds pages primarily by following links from pages it already crawls. Submitting a sitemap can speed discovery up, but the link graph is the mechanism. Placement is part of writing the page, not an afterthought once it is live.

REACHABLE	A few clicks from a page that is already crawled often — not stranded behind a search box, a filter, or an archive nobody links to.
LINKED	From the pages covering neighbouring intents, with anchor text that names the intent instead of saying “read more”. Write the anchor into the sentence, at drafting time.
LISTED	In the sitemap with an honest last-modified date: the date the content changed, not the date the build ran.
SINGULAR	Served at one canonical URL, with no second address serving the same content.

WHAT GOOGLE DOCUMENTS ABOUT SITEMAPS

ELEMENT	STATUS
Size limit	50 MB uncompressed, or 50,000 URLs, per sitemap file
<priority>, <changefreq>	Ignored
<lastmod>	Used if it is consistently and verifiably accurate — checked against the page's actual last modification
Significant update	Main content, structured data or links. A changed copyright year does not qualify

A site that stamps every URL with the deploy timestamp has told the engine its lastmod is worthless, and it will be treated accordingly.

PRIMARY SOURCES

- Build a sitemap
- SEO starter guide
- Crawl budget

robots.txt controls crawling. Not indexing.

Google states there are very few technical things a page must do, and most sites pass without noticing: Googlebot is not blocked; the page works, meaning it returns HTTP 200; and the page has indexable content. Meeting all three is eligibility, not a guarantee.

THE MOST EXPENSIVE SELF-INFLICTED BUG IN THE FIELD

A URL disallowed in robots.txt can still be indexed if other sites link to it — described only by those links. And for a noindex rule to work, the crawler must be able to fetch the page and read it. Put those together: a URL that is BOTH disallowed in robots.txt AND marked noindex can stay in the index indefinitely, because the directive that would remove it is behind the door that was locked. Allow crawling until the noindex has been seen and the URL has dropped out — only then, if you still want to, disallow it.

WHAT TO TEST, AND WHAT COUNTS AS PASSING

PRECONDITION	THE TEST	PASSES WHEN
Crawlable	Fetch as the crawler would; check the live robots.txt	No matching Disallow for that user agent
Indexable	Read the response headers and the rendered head	No noindex in either, unless intended
Returns 200	Follow redirects; record the final status	Final status is 200, and not a soft 404
One URL	Request http/https, www, trailing-slash and parameter variants	All variants redirect to one address
Distinct	Compare main content against the nearest siblings	Not a near-duplicate differing by a parameter
Canonical sane	Check the rel=canonical target	Self-referential, 200, indexable, not a redirect

The check log records what was actually tested, not what was assumed. NOT TESTED is an acceptable value. PASS without evidence is not. And do not reach for noindex to save crawl budget: Google will still request the URL and then drop the page, so the crawl is spent either way.


PRIMARY SOURCES

- robots.txt intro
- noindex
- Technical requirements

09 | RENDERING

Late is a cost, even when it is not invisible.

Google processes JavaScript in three phases. Googlebot fetches the URL and parses the HTML for links; the page is queued for rendering in a headless Chromium; the rendered HTML is indexed and its links extracted. Google says a page may sit on the rendering queue for a few seconds, but that it can take longer than that.



THE FIVE CHECKS

- Fetch the page without executing scripts and look for the main content, the title, the h1, the canonical and the robots directives. If they are not there, that is the finding.
- Check that every resource the renderer needs is crawlable. robots.txt applies to scripts and APIs too; block one and the renderer indexes an empty page.
- Use the History API for routing. Fragment routes are not distinct URLs to a crawler, so every view collapses into one indexed page.
- Make “not found” states return a real 404 or 410, or carry noindex. A single-page app that renders “not found” with HTTP 200 has produced a soft 404.
- Decide the rendering strategy in writing. Client-side-only is legitimate for an application UI; it is not legitimate for the page you are writing in this run.

DO NOT SOLVE THIS WITH A CRAWLER EXCEPTION

Pre-rendering the same content for everyone is fine. Detecting Googlebot and serving it a richer page is cloaking — presenting different content to users and search engines to manipulate rankings — and it is named in the spam policies. Server-side or pre-rendering is still, in Google's own words, a great idea: it makes the site faster for users and for crawlers alike.

PRIMARY SOURCES

[JavaScript SEO basics](#)
[Spam policies](#)
[Search Console](#)

Eligibility, not ranking.

Structured data does not improve rankings. It makes a page eligible for a matching rich result. That is the whole of the mechanism, and the honest way to describe it to a stakeholder. Google recommends JSON-LD as the easiest format to implement and maintain at scale.

THE TWO RULES THAT GET BROKEN

First: only mark up what the type actually describes — the markup must represent the page's real content. Second: never mark up content that is not visible to readers of the page, even if the information is accurate, and never create a page just to hold structured data. A structured data issue can result in a manual action, which removes rich result eligibility. The downside of over-marking is losing exactly the thing you were reaching for.

THE ANTI-PATTERNS, AND WHY EACH ONE FAILS

ANTI-PATTERN	WHY IT FAILS
Question-and-answer markup on questions nobody asked	Marks up content that exists only to carry the markup
Review markup for in-house reviews of your own product	Misleading markup; the review documentation constrains self-serving reviews
Product markup on a page with no product	Relevance failure — the markup does not represent the page
Markup for content behind a tab that never renders	Not visible to readers of the page
A competitor's JSON-LD with the strings swapped	Required properties silently wrong; it describes their page
Telling the board it will lift rankings	It is an eligibility mechanism, not a ranking one

“Article and Breadcrumb, and nothing else” is a complete, correct outcome of this skill. Google documents a finite gallery of feature types — among them Article, Breadcrumb, Carousel, Dataset, Discussion forum, Event, Job posting, Local business, Organization, Product, Profile page, Q&A, Recipe, Review snippet, Software app, Video. If none of them describes the page, none of them goes on the page.

PRIMARY SOURCES

- [Intro to structured data](#)
- [General guidelines](#)
- [Feature gallery](#)

A floor to clear, not a lever to pull.

Page experience signals are documented and real. Relevance dominates them. Google states there is no single page-experience signal, that Core Web Vitals are used by its ranking systems, that good scores do not guarantee a top position, and that Search always seeks to show the most relevant content even where the page experience is sub-par.

METRIC	WHAT IT MEASURES	GOOD	ASSESSED AT
LCP	Loading — render time of the largest image, text block or video in the viewport	≤ 2.5 s	75th percentile
INP	Interactivity — latency across click, tap and keyboard interactions in a visit	≤ 200 ms	75th percentile
CLS	Visual stability — the largest burst of unexpected layout shift over the page's life	≤ 0.1	75th percentile

Assessed separately for mobile and desktop. INP replaced First Input Delay as a Core Web Vital in March 2024 — a report still showing FID is stale. TTFB, FCP and TBT are useful diagnostics but are not Core Web Vitals.

DIAGNOSE IN THE LAB

Synthetic runs are reproducible and are the right tool for finding why something is slow. They are not the measurement anyone is scored on, and a green lab score proves nothing about real visits.

DECIDE ON THE FIELD

Search Console's Core Web Vitals report shows real-world usage data. It lags, because it is a rolling window over past visits — so a fix does not show up today, and that is not a reason to ship another one.

CHECK THE PRIMARY PAGE BEFORE ACTING ON A THRESHOLD

Third-party posts circulate claiming the LCP threshold tightened to 2.0 s, or that INP's methodology changed, during 2026. Google's own documentation and web.dev both still state 2.5 s and 200 ms. Soft navigations are moving toward measurement — enabled by default from Chrome 151, with the stated aim of inclusion once the API launches — but that is an aim, not a shipped Core Web Vital. This is exactly the documented-versus-believed distinction the playbook exists to enforce.

PRIMARY SOURCES

- Page experience
- Web Vitals
- Optimize INP
- Optimize LCP

There is no separate discipline. Yet.

Google states that its generative features on Search are rooted in its core Search ranking and quality systems, and that the best practices for SEO continue to be relevant because of it. Four of its statements are worth carrying into every meeting where somebody proposes an AI-specific workstream.

- There are no additional requirements to appear in AI Overviews or AI Mode, and no other special optimisations necessary.
- You do not need to create new machine-readable files, AI text files or markup to appear in these features, and there is no special schema.org structured data to add.
- The controls are the ordinary Search controls: robots.txt for Googlebot, and nosnippet, data-nosnippet, max-snippet or noindex to limit what is shown. There is no AI-only opt-out inside Search.
- Google-Extended is a separate token covering training and grounding in some of Google's other systems. Google documents that it does not affect inclusion in Search and is not a ranking signal — so it is not an AI Overviews opt-out.

WHERE BING IS DIFFERENT, AND WORTH READING

AREA	BING'S DOCUMENTED POSITION
Framing	The guidelines cover Bing search experiences, Copilot, and grounding API results together
AI directives	NOARCHIVE prevents use in Copilot responses and grounding; NOCACHE limits Copilot to URL, title and snippet
AI training	Bing documents noindex as also excluding content from training its generative foundation models
Rendering	Avoid hiding critical content behind client-side rendering; unrenderable content may not be indexed or grounded
Redirects	301 for permanent moves; 302 only for changes under about two days; prefer redirects to canonical tags
Discovery	IndexNow is listed first, ahead of sitemaps and internal links; streaming beats batch submission
Expectations	Bing states plainly that SEO does not guarantee rankings, and GEO does not guarantee citations

IndexNow is a ping protocol, not an indexing guarantee: participating engines include Microsoft Bing, Naver, Seznam.cz, Yandex and Yep — Google is not among them — and a 200 response only indicates the engine received the URL. Submitted URLs still count against crawl quota and are still subject to every quality gate. It accelerates discovery; it does not change eligibility.

PRIMARY SOURCES

- AI features
- Generative AI guide
- Bing guidelines
- IndexNow

13 | EVIDENCE DISCIPLINE

Tell your facts from your folklore.

Search ranking is not a published algorithm. Google names its ranking systems — BERT, RankBrain, neural matching, passage ranking, freshness, deduplication, link analysis, site diversity, spam detection — but not their weights or interactions, and they change without notice. Some retire outright: the helpful content system became part of core ranking in March 2024.

DOCUMENTED

The engine says so, in its own published documentation. Cite the URL. This is the only label that can carry a decision on its own.

OBSERVED

We saw it ourselves, on a named date, on a named URL, in a named locale. It decays, so it is re-checked rather than remembered.

BELIEVED

Widely held in the field, not confirmed by any engine and not tested by us. Never the basis for a decision on its own.

TWO RULES THAT FOLLOW

- No claim that a specific change causes a specific ranking movement — the input space is not observable from outside.
- No treating a correlation study as a mechanism. It describes what ranking pages have in common, not what put them there.

CLAIMS THIS PLAYBOOK DELIBERATELY EXCLUDES

Recorded so nobody re-adds them. That the LCP threshold tightened to 2.0 s in 2026, or that INP's methodology was tightened — both contradicted by the primary documentation. That Google reads llms.txt — its AI features page does not name the file, and states no new machine-readable files are needed. That Google-Extended removes you from AI Overviews — Google documents it as covering some of its other systems. And any percentage, multiplier or ranking-factor weight, because no engine publishes one.

A team that cannot tell its documented facts from its folklore will spend a quarter optimising something that never mattered, and it will not find out.

PRIMARY SOURCES

[Ranking systems](#)
[AI features](#)
[Web Vitals](#)

The cheapest output, and often the right one.

Refusal is a first-class output. The skill stops and reports — it does not publish — when any of five conditions is true. The report names what would have to be true first.

01

NOTHING FIRST-PARTY

The page would repeat what is already widely published, with nothing of the organisation's own added.

02

ALREADY SERVED

The intent is already served by an existing page on the same site, and no merge decision was taken.

03

UNSOURCEABLE CLAIM

A central claim needs a number the organisation cannot source — and inventing it is not an option.

04

FORMAT WE CANNOT PRODUCE

The winning format is one the site cannot credibly produce, such as independent third-party evaluation.

05

THE PAGE EXISTS TO HOLD A PHRASE

It was proposed because a competitor has that URL, not because a reader has that question. This is the condition that produces forty pages by March if it is not named the first time.

THE ECONOMICS, FOR THE CONVERSATION THAT FOLLOWS

PUBLISH IT ANYWAY	STOP AND REPORT
It competes with your own better pages	The skill reports instead of publishing
It dilutes the internal links into that section	The report names what would have to be true first
Somebody has to maintain it who did not want it	No URL is added for the section to carry

Conditions 1 and 5 are not only editorial judgements. Google's spam policies define scaled content abuse as many pages generated primarily to manipulate rankings rather than help users, “no matter how it's created”, and name generative tools producing many pages without adding value as an example. Refusing is cheap. The alternative is not.

PRIMARY SOURCES

- Spam policies
- Generative AI content
- Helpful content

Query drift is the signal that pays.

The measurement that exists is your own. Search Console gives you the queries a page attracted, how often it was shown, where it sat and what was clicked — first-party ground truth rather than a third-party estimate. Set the review date before publishing, along with what you would change if the result disappoints, so the review is a decision and not a browsing session.

READ IT IN THIS ORDER	WHAT IT MEANS
Is it indexed at all?	Everything below is meaningless until this is yes
Impressions on seed queries	The engine is considering the page for the intent
Query drift	The promise does not match the page — the most useful signal on this list
Average position band	Movement, not a number to report to two decimal places
Clicks vs. definition of done	The only measure of whether the page did its job
Rich results and manual actions	Eligibility retained or lost; a policy problem, not a ranking one

WHEN A PAGE COLLECTS THE WRONG QUERIES, FIX THE PROMISE FIRST

When a page starts collecting queries other than the intent it was written for, the problem is usually the promise rather than the word count: the title, the opening and the headings are describing a page somebody else wrote. Fix the promise. Wait for a full re-crawl. Only then ask whether the page needed to exist. Do not reorder those three — rewriting the body first changes the thing that was probably not broken.

EVERY REVIEW ENDS IN ONE OF FOUR WRITTEN DECISIONS

- KEEP — it is doing its job. Set the next review date and leave it alone.
- FIX THE PROMISE — drift detected. Rewrite title, opening and headings only.
- IMPROVE THE SUBSTANCE — the promise is right and the answer is thin. Return to the page contract.
- RETIRE OR MERGE — the page should not have existed. Redirect it and take the lesson into the next brief.

On timing, Google's own framing: some changes take effect in a few hours, others take months, and in general it is worth waiting a few weeks before assessing whether the work helped. Note too that AI Overviews and AI Mode traffic is included in overall Search traffic in Search Console rather than broken out — anyone reporting an exact “AI traffic” figure from it is reporting something the tool does not provide.

PRIMARY SOURCES

- Search Console
- SEO starter guide
- AI features

An incomplete set is a failed run.

A run is complete only when all five exist. The one exception is a refusal, which is complete with the brief, the decision and the refusal report. Templates for every one of them ship in the companion skill pack, under `references/output-templates.md`.

THE BRIEF	Intent, seed queries, the observed shape of the current results, and the date that observation was made.
THE DECISION	New page, improve, merge or decline — with the reasoning and the URLs affected.
THE DRAFT	The page itself, plus the title element and meta description, written against the page contract.
PLACEMENT PLAN	Where the page sits, which existing pages link into it and with what anchor text, and which links it sends out.
THE CHECK LOG	Each precondition, whether it passed, and what was actually tested rather than assumed.

WHAT SHIPS IN THE COMPANION PACK

FILE	WHAT IT IS
SKILL.md	The orchestrator: order of work, handoffs, output contract, refusal rules
skills/ × 10	One folder per skill — procedure, checklist, copy-ready prompt, output template, stop conditions, sources
references/references.md	Every source URL, what it contributed, and the claims deliberately excluded
references/output-templates.md	The five artefacts and the refusal report, ready to fill
references/crawlers-and-other-engines.md	Google's three crawler families, the Bing guidelines, and IndexNow
examples/ × 2	A full run that ships, and a full run that refuses

DOWNLOAD THE COMPANION SKILL PACK

Ten skills, two worked examples and the source table, as Markdown you can drop into `~/.claude/skills/`, into a repository under `.claude/skills/`, or into a Claude Project as project knowledge. Free.

get-tau.com/resources/seo-skill-pack-skills.zip →

Check the primary page before you act on it.

Every claim in this playbook traces to one of the pages below, all read on 3 September 2026. Search documentation changes without notice. Before acting on a threshold, a limit or a policy definition quoted here, open the primary URL and confirm it still says what this page says it says. That habit is the point of the playbook, not a disclaimer on it. No statistic in this document is invented, and none is presented without a source.

GOOGLE SEARCH CENTRAL — FUNDAMENTALS

[Search Essentials](#)
[Technical requirements](#)
[Spam policies](#)
[SEO starter guide](#)
[Creating helpful content](#)

CRAWLING AND INDEXING

[robots.txt introduction](#)
[Block indexing with noindex](#)
[Consolidate duplicate URLs](#)
[Build and submit a sitemap](#)
[Crawl budget for large sites](#)
[JavaScript SEO basics](#)
[Google crawlers overview](#)

APPEARANCE IN RESULTS

[Title links](#)
[Snippets](#)
[Intro to structured data](#)
[Structured data guidelines](#)
[Feature gallery](#)
[Page experience](#)
[Ranking systems guide](#)
[AI features and your website](#)

GENERATIVE AI. MONITORING. PERFORMANCE. OTHER ENGINES

[Optimizing for generative AI](#)
[Using generative AI content](#)
[Search Console overview](#)
[Web Vitals](#)
[INP](#)
[LCP](#)
[CLS](#)
[Optimize INP](#)
[Optimize LCP](#)
[Soft navigations](#)
[Bing Webmaster Guidelines](#)
[IndexNow documentation](#)

READ THE ARTICLE THIS PLAYBOOK CAME FROM

The SEO skill spec sets out the same method as prose: the inputs the skill needs, how it reads a result page, the contract a published page has to satisfy, and the conditions under which it refuses to publish.

get-tau.com/resources/seo-skill-spec →

Free diagnostic — see where your own site stands across search, AI and social before you run any of this:
platform.get-tau.com/diagnostics