


TAU FIELD GUIDE | SEPTEMBER 2026

LISTENING THAT ENDS IN A REPLY

Finding the public conversations where a reply would genuinely help somebody — and where the product is genuinely part of the help.

8 skills | SOURCES → INTENT → SCORE → VETO → DEDUPE → DRAFT → REVIEW



A listening system earns its place the moment it can stay silent. Everything here that looks like friction exists because the cost of a bad reply is paid by the brand, in public, permanently.

TAU

Marketing autopilot for new products.

00 | HOW TO USE THIS GUIDE

The skill drafts. A person posts.

This guide is the reasoning behind the eight skills in the accompanying pack. Each section states a decision, the rule that governs it, and the evidence the rule rests on. Where a claim comes from a platform policy or a regulation, the source is linked at the foot of the section and listed in full on the last page.

01 Sources and access	Which communities are in scope, through access the platform actually grants.
02 What each platform allows	The access surface, automation rules and promotion rules of seven platforms.
03 The retrieval net	Four term sets, homonym rules, dated exclusions, and a measured precision figure.
04 Intent	The quoted evidence sentence, and the seven tests keyword matching fails.
05 Early signal	What is actually known about catching a thread while it moves.
06 Scoring and the veto list	An order for the queue, and — kept strictly apart — what vetoes an item.
07 Dedupe and cadence	One conversation, one reply. Caps, cooldowns and expiry.
08 The reply	Answer first, disclose, one link at most, concede the limits, no hook.
09 Escalation	The five kinds of draft a person has to own before anything is posted.
10 The standing review	Reading the replies rather than the counts, and the five failure modes.
11 Conduct and privacy	What a listening system may read and keep, and the output contract.
12 Disclosure	The one rule that does not depend on the platform.
13 Numbers	What this guide refuses to tell you, and why.
14 The pack	The eight skills, the references file, and the two worked examples.
15 Sources	Everything the guide rests on, linked.

THE ONE RULE

The skill drafts and a person posts. That division is where the whole design stays defensible, and it holds until the queue has a long record of a human agreeing with the draft — not until the queue is merely large.

- A reviewer's only two options are approve-and-post, or reject with a reason.
- A reviewer who rewrites is telling you the drafting step is not finished. Log it as a rejection.
- A queue with no rejection record cannot be reviewed later.

01 | SOURCES AND ACCESS

Only public posts, only through access the platform grants.

That is a boundary rather than a preference. Terms of service, rate limits and the difference between an official interface and scraping around one are part of the specification, not an operational detail to sort out later. Two constraints apply to every source: anything requiring evasion of a rate limit, an authentication wall or a robots directive is out of scope, and anything a platform prohibits does not become acceptable because an automated system is doing it.

01

DISCUSSION FORUMS

Threaded communities where questions are asked and answered, and where local rules on self-promotion run from tolerated to banned. Read each community's rules before adding it, and store them with the source.

02

SHORT-FORM SOCIAL

Fast, high-volume, low-context. Strong for spotting something starting to move, weak for judging intent from the text alone.

03

PROFESSIONAL NETWORKS

Slower, identity-attached, more permanent. A reply here is closer to a public statement by a named person than to a comment, which raises both the value and the cost.

04

COMMENT SECTIONS

Video and article comments, where questions accumulate under an evergreen piece long after publication and often stay unanswered.

05

CLOSED COMMUNITIES

Chat servers and private groups. Include only where participation is invited and the operator knows what the account is, and never treat a member list as an audience to be harvested.

06

WHAT IS NOT A SOURCE

Anything reachable only by scripting a website, masking a user agent, or working around an auth wall. Record the exclusion rather than leaving the source unlisted.

THE ACCESS ROUTE IS PART OF THE SOURCE RECORD

Three of the seven platforms surveyed for this guide changed their access terms between late 2025 and early 2026. A source map without a date on each entry is not a record of permission; it is a record of what somebody believed once.

- Store the platform rules and the community rules separately — each with a URL and the date read.
- Treat an entry older than your review cadence as unread, not as permission.
- Where a platform prohibits an access route, record the decision, not just the absence.

02 | WHAT EACH PLATFORM ACTUALLY ALLOWS

Read at source on 3 September 2026.

Every line below was read from the platform's own documentation or policy page on that date, and every URL is on the closing sources page. Re-read anything you are about to depend on: this table will go stale, and the guide would rather say so than pretend otherwise.

Hacker News	Open API, no authentication, and the documentation states there is currently no rate limit. Full-text search is available separately. But the site guidelines ban generated or AI-edited text outright: "HN is for conversation between humans." A source to listen to; not a destination for a drafted reply.
Stack Overflow	API with a published 10,000-unit daily quota, a 30-requests-per-second IP throttle, and mandatory attribution. Product teams are explicitly told to monitor their tags — and told, in as many words, that if you mention your product you must disclose your affiliation in the post.
YouTube comments	Comment reads cost one quota unit each, but the daily allocation is not a flat ten thousand calls: search is capped separately at 100 calls a day, and non-authorised data may be stored for at most 30 days. Enumerate comments on known videos rather than burning search.
X	Brand-mention monitoring is listed as an approved use case, priced per resource read since the move to pay-per-use in February 2026. Auto-replying to anyone mentioning a keyword is prohibited, and an app that generates and posts replies needs prior approval. Non-API automation risks permanent suspension.
Reddit	Access now requires explicit approval rather than self-service registration, OAuth is mandatory, and commercial use requires a contract. Deletions are expected to propagate to your store. The operative rule on promotion is the individual community's, not any site-wide ratio.
LinkedIn	There is no listening API: the permission to read member posts is closed. Scraping is prohibited by the User Agreement and by an explicit notice in robots.txt. A place a person participates, not a place a system reads at scale.
Discord	A bot sees only the servers it was invited to, and reading message text needs a privileged intent. The Developer Policy prohibits mining, scraping, profiling users, and contacting them without explicit permission.

THE "9:1 RULE" IS NOT A PLATFORM RULE

The most repeated number in social listening does not appear in any enforceable Reddit policy. Reddit's content policy states the rule without a ratio: participate authentically in communities where you have a personal interest, and do not spam. The ratio survives on a page that disclaims its own status as an informal expression of values, and in guidance describing what some communities choose to adopt.

- Which is why this method records a participation state per community, from that community's own rules, rather than applying one global ratio everywhere and calling it compliance.

03 | THE RETRIEVAL NET

Cast wide on purpose, then read everything it caught.

Keyword matching is the retrieval layer, not the decision. Its job is recall: it casts a wide net over product names, category terms, problem descriptions and the phrasings people use when they do not know the category term. Precision is bought back in the next stage, by reading. A query tuned for precision at the retrieval stage buys it by throwing away exactly the posts written by people who have not yet learned your vocabulary — which is most of the high-intent ones.

SET A**BRAND TERMS**

The product name and its misspellings. Needs a homonym rule whenever the name is also an ordinary word, a place, a person or another company's brand.

SET B**CATEGORY TERMS**

What the thing is called by people who already know the category. High precision, and systematically blind to newcomers.

SET C**PROBLEM PHRASINGS**

How people describe the pain before they know the category term. The highest-intent set, and the one most often missing entirely.

SET D**COMPETITOR TERMS**

Names of alternatives. Surfaces switching intent, and simultaneously triggers a veto when the competitor is the thread's subject.

WHERE SET C COMES FROM

Mine it from language people actually used: support tickets, sales-call notes, review text, and existing threads. Do not build it from the product's own marketing copy — that vocabulary is precisely the one prospects do not have yet.

Intent modifiers

Keep them as a separate dimension rather than baking them into terms: "recommend", "alternative to", "how do I", "anyone using", "worth it", "switching from", "tired of". Combined with sets B and C they raise precision without cutting recall on the term itself.

Exclusions

Build the list from what the query actually caught, not from guesswork. Run it, read a sample, add an exclusion for each recurring wrong-topic match, and date it so it can be revisited when the topic moves on.

Per-source syntax

Operators, boolean support and result ordering differ by platform, and a query that behaves on one source will silently behave differently on another. Store the query per source, never once globally.

Versioning

A query is a configuration with an owner and a date. Any change to it invalidates the precision figure measured on the previous version.

04 | INTENT

A classification with a reason, not a similarity score.

Retrieval matched a string. Triage decides whether a person in that thread actually wants something. Every queue item records the sentence that established intent, quoted verbatim. If none can be quoted, the item is dropped — and a drop is a success.

Intent	Is there a job to be done? The strong forms are an explicit request for a recommendation, a problem with no solution yet, frustration with a current approach, and a stated plan to change something. Everything else is discussion, and discussion is not intent.
Fit	Would this person plausibly be helped by what you have? Role, scale, stack and vocabulary all carry it. A thread full of practitioners describing a problem you do not solve is a bad fit however relevant the keywords look.
Timing	Is the conversation live? Recency, whether replies are still arriving, whether the question is answered, and whether the thread is locked or archived. A perfect answer to a settled question is noise.
Reach	How many will read this later? Threads that rank in search keep receiving readers for years, which can matter more than the audience present today. Record it as an observation, never an estimated number.
Risk	What could a reply cost? Community rules, a competitor thread, a complaint about your own product, a regulated context, personal distress, a pile-on. Risk does not lower a score — it vetoes.

THE SEVEN TESTS A KEYWORD-ONLY SYSTEM FAILS

Each of these needs its own test, run explicitly, with its outcome recorded even when it passes. They are the reason the full thread has to be read rather than the matched snippet.

- **Negation** — the post says it is not looking for a tool, or has already chosen one.
- **Past tense** — the problem was solved last month and the thread is a retrospective.
- **Sarcasm and complaint** — the category term appears inside criticism rather than a request.
- **Homonyms** — a product name that is also an ordinary word, a place, or somebody else's brand.
- **Quotation** — the post describes what someone else asked rather than asking.
- **Answered questions** — a reply already solved it and another voice adds nothing.
- **Language and register** — a machine translation reads as a stranger arriving in a local community.

COPY-READY TRIAGE PROMPT

Read the full thread, not the matched snippet. Quote verbatim the one sentence that establishes a job to be done; if you cannot quote one, drop the item as no-evidence. Assign exactly one intent class, run all seven rejection tests and record every outcome including the passes, and assess fit against what the product actually does.

- Never infer intent from the keyword that matched.
- Never estimate an audience figure you cannot observe on the page.
- A triage step that never drops anything is not reading.

05 | EARLY SIGNAL

What is actually known about catching a thread while it moves.

Timing is the dimension most often asserted and least often evidenced. There is real published work here, so this section cites it — with its dates, because the foundational studies are a decade old and the platforms they measured have changed.

Growth is predictable, relatively

Cheng et al. (WWW 2014) reframed the question from “how big will this get” to “given a cascade of size k , will it double?” — a formulation where guessing scores 50 per cent. They reported 79.5 per cent accuracy, rising as more of the cascade is observed, and found **temporal features the single strongest class**: cascades whose early reshares come quickly are the ones that grow. Measured on Facebook photo reshares in 2014.

One hour of observation

Zhao et al. (KDD 2015) reported roughly 15 per cent relative error in predicting a cascade’s final size after observing it for one hour, using only timestamps and graph degrees. The operationally useful output was not the size but the state: is this past its phase of explosive growth?

The content ceiling

Dogan et al. (2025–26), on 46,578 Reddit posts across 25 communities, found content-only models plateau: structural and temporal features are needed to pass the limit. Their precision-recall AUC ran from 0.43 at 30 minutes to 0.80 at seven hours. **Reading the post text alone will not tell you whether a thread is moving.**

What the early signal is made of

The same work describes an evidentiary transition: at 30 minutes the prediction rests on who posted and where; by seven hours it rests on velocity and acceleration. Early on you are reading the author and the community, not the curve.

Normalise by community

That study defines virality relative to community size and engagement dynamics rather than by an absolute volume threshold, criticising the field’s reliance on “arbitrary cut-offs”. A freshness window and a reach note should be per-source for the same reason.

WHY ARRIVING EARLY IS NOT THE SAME AS ARRIVING FAST

In a randomised experiment on 101,281 comments (Muchnik, Aral and Taylor, Science, 2013), a single artificial early up-vote raised the probability that the next viewer up-voted by 32 per cent, and lifted final mean ratings by 25 per cent over five months. Negative manipulation was corrected by the crowd; the effect was asymmetric.

- The finding is about early engagement, not early posting: what arrives first shapes what follows.
- It argues for being in the thread while it is live — and against manufacturing that engagement yourself, which every platform’s manipulation policy prohibits.

06 | SCORING AND THE VETO LIST

Two mechanisms that must never be one.

Scoring exists to order a queue that is longer than the time available. It does not decide whether to reply. That decision belongs to the veto list, and the two stay apart — because the moment risk becomes a score penalty, a high enough intent can argue its way past a rule.

SCORING

ORDERS A QUEUE

Intent and fit dominate the shape. Timing modulates what they produced rather than creating relevance. Reach breaks ties only. Risk is not in the arithmetic at all.

VETO

DECIDES ADMISSION

Checked on its own, against the stored rules. A vetoed item never reaches the queue, whatever it scored. Nothing an item scores can argue past a veto.

THE VETOES

- The community prohibits promotional replies, or prohibits links, and the reply would breach that.
- The thread is a complaint about your product, which belongs to support and not to outreach.
- The thread names a competitor as its subject, where a reply from you is an advert wearing a helpful voice.
- The context is legal, medical, financial or safety-critical, or the author is in evident distress.
- The thread is a pile-on, a brigade or an active controversy.
- The account has already replied in this thread, or to this author recently, or has hit the community's cap.
- The post is older than the freshness window agreed for that source.
- The reply would need an undisclosed or second account — one identity, disclosed, or nothing.

WEIGHTS ARE AN OPINION WITH A VERSION NUMBER

Any weights you start with are guesses. They are worth writing down so they can be argued with and adjusted against what actually happens, and they are not evidence. Nothing in this guide supplies a weight, because a weight that has not been measured on your own queue would be a fabricated number dressed as a default.

- Cut the queue to the time budget, not to a score threshold.
- If the queue is always longer than the budget, narrow the query — do not lower the bar.
- Re-fit weights against outcomes: survival, moderator action, author response. Not engagement.

07 | DEDUPE AND CADENCE

The same question travels. Three replies read as a campaign.

One person asks in a forum, links it on social, and someone else quotes it a day later. Three queue items for one conversation produce three replies from the same account, and that is the single most reliable way to get noticed for the wrong reason.

Group by conversation	Not by post. Match on the author across platforms, the shared link or quoted text, and the substance of the question — and require more than one signal before grouping.
One carrier, one reply	Choose the item in the most permissive community with the liveliest thread to carry the reply, mark the rest duplicate-of with a reason, and let a person decide if a second is ever warranted.
Cap per community	Counting queued drafts as well as posted replies. A cap enforced only at posting time is not a cap. An account that shows up in every thread becomes the thing people notice.
Cool down per author	Across communities, not within one. Following one person from place to place reads as surveillance regardless of intent, and is one of the few behaviours that reliably produces a complaint.
Expire, do not re-rank	Items not acted on inside their freshness window are dropped, never carried forward. A stale reply is worse than silence, and an ageing queue is how a bar quietly drops.
Honour deletion	If the source post has been deleted or edited away by the time a person reaches it, drop the item. Never reply from a cached copy of something someone took down.

RE-CHECK AT POSTING TIME

A conversation that was one thread when queued may be three by the time a person gets to it, and the duplicate may already have been answered — by someone else, or better. The grouping is re-run immediately before posting, not only when the queue was built.

08 | THE REPLY

Would this help if every mention of your product were deleted?

The test is simple and hard to pass. Apply it literally: delete every product mention from the draft and read what is left. If what remains is empty or generic, the draft is an advert, and the community will read it as one.

Answer first	Open by addressing the actual question, in enough detail to be usable on its own. A reply whose value is behind a link has given nothing.
Disclose the connection	Say who you are and that you work on the thing you are mentioning, in the reply itself rather than in a profile. It is the rule in many communities, and undisclosed endorsement by someone with a material connection to the product is a regulatory problem in several jurisdictions.
One link at most	And only when the link is the answer. A reply carrying two links is being written for the person posting it. Where the community restricts links, carry none.
Match the register	Communities have a voice. Read three recent well-received replies before drafting. Marketing phrasing in a technical thread is detected instantly and read as an intrusion, whatever the content underneath.
Concede the limits	Say plainly where your product does not fit and what you would use instead. It reads as honesty because it is honest, and it costs a sale you were not going to make.
Leave without a hook	No closing pitch, no offer, no request to book anything. The reply ends when the answer ends.
Sometimes, do not mention it	A helpful answer with no product in it is a legitimate and complete output. A skill that cannot produce one will eventually force the product into a thread where it does not belong.

THE DRAFT IS POSTED UNEDITED OR REJECTED

A draft that needs a rewrite by the reviewer has saved nobody any time. Attach to every draft the disclosure line, the link count, the register note, and the result of the deletion test — so the reviewer is approving evidence, not prose.

- Claim nothing you could not support if a person wrote it in public copy.
- No performance figures, no comparisons, no “fastest” and no “only”.
- Length is not thoroughness. A reply that answers in a paragraph beats one that answers in six.

09 | ESCALATION

Five drafts a person has to own.

The skill drafts. A person posts. Within that, five kinds of draft need more than a reviewer's glance, because the accuracy bar is the one you would accept in public copy rather than the one you would accept in a comment.

HAND OFF**YOUR OWN PRODUCT'S FAILURES**

When the thread touches something your product got wrong, whoever owns support takes it, with the thread attached. Do not post a holding reply in the meantime.

HAND OFF**A NAMED COMPETITOR**

When the reply would compare you to a competitor by name, a person writes it, at the accuracy bar you would accept in public copy.

REVIEW**A REGULATED CONTEXT**

Where claims are regulated, whoever signs those claims off reviews the draft — with the evidence and the veto record, not only the text.

REVIEW**CORRECTING THE AUTHOR**

When the reply would tell the author they are wrong, a person reviews it. Being right in public is not automatically worth doing.

STOP**A MODERATOR REMOVAL**

Once a moderator has removed the account from a community, that source pauses until a person re-reads its rules and re-dates the entry.

STANDARD**EVERYTHING ELSE**

The named reviewer approves and posts, or rejects with a reason. A reviewer who rewrites is telling you the drafting step is unfinished; log it as a rejection.

10 | THE STANDING REVIEW

Read the replies, not the counts.

A listening system that cannot see what happened to its replies is running blind, and it will keep making the same mistake in a new community every week. The feedback that matters is not engagement. It is whether the reply survived, whether a moderator removed it, whether the author responded and how, and whether anyone else in the thread objected to its presence.

THE FIVE FAILURE MODES, IN ROUGHLY THE ORDER THEY APPEAR

- **Volume creep** — the queue grows and the bar quietly drops to fill the time available.
- **Keyword drift** — a broad term starts pulling in a neighbouring topic nobody noticed.
- **Voice collapse** — every reply reads as the same paragraph with the nouns swapped.
- **Rule rot** — a community changed its policy and the stored copy did not.
- **Removal blindness** — nobody checks which replies were taken down, or the pattern in them.

READ THE REMOVALS AS A SET

One removal is an incident. Three removals citing the same rule is a source-map error, and the fix belongs in the source map rather than in the next draft.

THE REVIEW QUESTION

It is not how many conversations were joined. It is whether a regular in each of those communities would still be glad the account is there. If the answer is no for a community, pause it.

- Route findings by name: query changes to the query skill, rule changes to the source map, weight changes to scoring, register notes to drafting — each with the evidence that prompted it.
- A review that changes nothing and records nothing did not happen.

11 | CONDUCT AND PRIVACY

A listening system reads people who did not ask to be read by a company.

The constraints follow from that, and they are not negotiable against volume. They are also the constraints that keep the programme defensible if anyone ever asks what it collects and why.

- Public posts only. Never private messages, and never a private group joined under a pretext.
- One identity, disclosed, rather than several accounts creating the appearance of independent voices.
- No profiles of individuals beyond what the current conversation requires, and no linking of accounts across platforms to build one.
- Deleted and edited posts are honoured: a queue item whose source has gone is dropped, not acted on from a cached copy.
- Retention has a limit and a reason, because a permanent archive of what strangers said about your category is a liability nobody asked for.
- Anything requiring evasion of a rate limit, an authentication wall or a robots directive is out of scope, and stays out of scope when an automated system is doing it.

THE OUTPUT CONTRACT

Each queue item carries the same fields, and a missing field blocks the item rather than degrading it.

01

SOURCE AND LINK

Platform, community, permalink, and the timestamp the item was captured.

02

THE EVIDENCE

The quoted sentence establishing intent, plus enough thread to judge whether it is already answered.

03

CLASSIFICATION

Intent class, fit, timing state and reach note — each with the reason, not only the value.

04

RISK FLAGS

Every veto checked, with the community rules that applied and where and when they were read.

05

THE DRAFT

The proposed reply with its disclosure line, written to be posted unedited or rejected outright.

06

EXPIRY

The time after which this item is dropped rather than posted late.

12 | DISCLOSURE

The one rule that does not depend on the platform.

Community rules vary. This does not. Where someone with a material connection to a product recommends it in public, the connection has to be disclosed — and the regulator has published an example that describes this exact fact pattern rather than an analogous one.

FTC ENDORSEMENT GUIDES, 16 CFR § 255.5(b)(8)

“An online community has a section dedicated to discussions of robotic products... Unbeknownst to this community, an employee of a leading home robot manufacturer has been posting messages on the discussion board promoting the manufacturer’s new product. Knowledge of this poster’s employment likely would affect the weight or credibility of the endorsements. Therefore, the poster should clearly and conspicuously disclose their relationship to the manufacturer.”

- The same example places duties on the employer: train employees, and where you direct or have reason to know of such posts, monitor them.

Material connection	§ 255.5(a): a connection that “might materially affect the weight or credibility of the endorsement” and is not reasonably expected by the audience must be disclosed. It “can include a business, family, or personal relationship” — employment is squarely inside it, and no payment is required for it to bite.
Clear and conspicuous	§ 255.0(f): in an interactive medium such as social media, the disclosure “should be unavoidable.” The 2024 reviews rule is stricter still — 16 CFR § 465.1(c)(4): a disclosure “is not clear and conspicuous if a consumer must take any action, such as clicking on a hyperlink or hovering over an icon, to see it.”
Why not the profile	That sentence is the whole argument for putting the disclosure in the reply body. A bio is a click away, and a click away is, by the rule’s own words, not clear and conspicuous.
Insider reviews	16 CFR Part 465, effective 21 October 2024, makes it a rule violation — not merely a deceptive practice interpretation — for an officer or manager to write an undisclosed review, and reaches the business for disseminating an employee’s undisclosed testimonial.
United Kingdom	DMCC Act 2024, Schedule 20 ¶ 25 prohibits “falsely representing oneself as a consumer”; ¶ 13 covers concealed-incentive reviews. The CMA’s guidance says plainly that a reviewer with a financial interest — “for example is an employee or shareholder of the trader” — is information whose absence may mislead.
European Union	Unfair Commercial Practices Directive, Annex I point 22: falsely creating the impression that the trader is not acting for business purposes, or falsely representing oneself as a consumer, is unfair in all circumstances.

13 | NUMBERS

What this guide refuses to tell you.

A number without a source is not a default. It is a fabrication with a confident tone, and in a document about not making claims you cannot support it would be an odd thing to ship. Everything below is widely repeated, was searched for during this research, and could not be traced to a primary source — so it appears here instead of in the method.

- **Reddit API pricing.** Per-call and per-month figures circulate widely. Reddit publishes no commercial rate card; its own wording is only that some developers “may be charged fees.”
- **X subscription tier prices.** Only the per-unit pay-per-use prices could be confirmed from official documentation. The commonly quoted monthly figures could not.
- **A rate limit for the Hacker News search API.** The circulating figure traces to a third-party package README, not to the provider. Back off on a 429 rather than trusting a number.
- **Maximum penalties** under the FTC reviews rule, and the frequently quoted “10 per cent of global turnover” for the UK regulator — neither was found in a primary source in this pass.
- **Any response-time or conversion benchmark for brand replies.** Numbers of this kind circulate without a traceable sample, method or date. This guide quotes none.

THIS APPLIES TO YOUR OWN CONFIGURATION TOO

The scoring weights, the freshness windows and the cadence caps are the same kind of number. Nothing in this guide supplies one, because a weight that has not been measured on your own queue would be exactly the thing this page is about.

- Write your weights down so they can be argued with. Version them. Name their owner.
- Measure precision on a stated sample size, and date it.
- Re-fit against what actually happened to your replies — survival, moderator action, author response.

14 | THE PACK

Eight skills, a references file, and two worked examples.

The accompanying zip contains one folder per skill, each with a SKILL.md carrying frontmatter, a procedure, a checklist, example prompts and an output template. Drop the folders into `~/.claude/skills/` so each lives at `~/.claude/skills/<name>/SKILL.md`, or upload the files to a Claude Project as project knowledge and name the skill you want.

01

listening-source-map

The watch list: access routes, platform and community rules recorded separately with dates, a participation state per community, freshness windows and cadence caps.

02

listening-query-design

Four term sets, homonym rules, dated exclusions, and a precision figure measured on a stated sample rather than felt.

03

intent-triage

The quoted evidence sentence, one intent class, and the seven rejection tests run explicitly with their outcomes recorded.

04

candidate-scoring

An order for a queue longer than the time available, with the components visible and risk kept out of the arithmetic.

05

veto-and-risk-check

The absolute list, checked separately from the score, with the rule text and where it was read attached to every check.

06

conversation-dedupe

One conversation, one reply. Caps counted across queued and posted, per-author cooldowns, expiry, and deletions honoured.

07

reply-drafting

The deletion test, the disclosure line, one link at most, register matched to three read replies, a conceded limit, no hook — and the no-mention draft.

08

escalation-and-review

Five triggers routing drafts to the person who must own them, then the standing review that reads replies rather than counts.

references.md

Every source behind the pack, with its URL, what it supports, and a closing list of the claims that could **not** be verified and must not be re-introduced.

examples/

Two worked examples end to end. The first is a candidate that survives every gate and produces a disclosed, no-link reply. The second is a batch where the highest-scoring item is vetoed outright and the one reply that goes out mentions no product at all.

15 | SOURCES

Everything this guide rests on.

Read at source on 3 September 2026. The pack's **references.md** carries the same list with the quoted text of each rule and a closing section naming the claims that could not be verified.

DISCLOSURE: UNITED STATES

[FTC Endorsement Guides, 16 CFR Part 255](#)
[Endorsement Guides, Federal Register \(88 FR 48092\)](#)
[FTC: what people are asking](#)
[Consumer Reviews and Testimonials Rule, 16 CFR Part 465](#)
[Part 465, Federal Register](#)
[FTC: reviews rule questions and answers](#)

DISCLOSURE: UNITED KINGDOM AND EUROPEAN UNION

[DMCC Act 2024, Schedule 20](#)
[CMA208, Fake reviews guidance](#)
[UCPD 2005/29/EC, Annex I](#)

REDDIT

[Responsible Builder Policy](#)
[Data API wiki](#)
[Accessing Reddit data](#)
[Content policy](#)
[Data API terms](#)
[Spam policy](#)
[Keeping spam out of your community](#)
[Reddiquette](#)
[robots.txt](#)

X

[API pricing](#)
[Developer guidelines](#)
[Rate limits](#)
[Developer policy](#)
[Pay-per-use announcement](#)
[Automation rules](#)
[Authenticity policy](#)

LINKEDIN

[User agreement](#)
[Professional community policies](#)
[Crawling terms](#)
[Community Management API](#)
[robots.txt](#)

HACKER NEWS

[Guidelines](#)
[FAQ](#)
[Show HN](#)
[Official API](#)
[Search API](#)
[robots.txt](#)

EARLY SIGNAL AND INTENT: PUBLISHED RESEARCH

[Cheng et al., Can Cascades be Predicted? \(WWW 2014\)](#)
[Zhao et al., SEISMIC \(KDD 2015\)](#)
[Dogan et al., Early prediction of meme virality on Reddit](#)
[Muchnik, Aral & Taylor, Social Influence Bias \(Science 2013\)](#)
[Gupta et al., Identifying Purchase Intent from Social Posts \(ICWSM 2014\)](#)
[Amigo et al., WePS-3 online reputation task \(CLEF 2010\)](#)
[Herrera-Poyatos et al., LLM sentiment uncertainty \(2025\)](#)

AI TRANSPARENCY AND QUERY METHOD

[EU AI Act, Article 50](#)
[EC Code of Practice on AI-generated content](#)
[Talkwalker, boolean operators](#)
[Brandwatch, writing a good query](#)
[Meltwater, boolean media monitoring](#)
[ESOMAR social media research guideline](#)

DISCORD, STACK OVERFLOW, YOUTUBE

[Discord developer policy](#)
[Discord rate limits](#)
[Discord gateway intents](#)
[Discord community guidelines](#)
[Stack Overflow: promotion](#)
[Stack Overflow: product support](#)
[Stack Exchange API throttles](#)
[Stack Exchange API terms](#)
[YouTube commentThreads.list](#)
[YouTube quota costs](#)
[YouTube developer policies](#)
[YouTube spam policy](#)