FEASIBILITY MAP · OCTOBER 2026

Building social listening with Claude Code

What you can build, whether it's worth building, and what you can't build at all, layer by layer.

29 / 33
features fully buildable

4
features only partly buildable

~30–50%
of a commercial tool's source coverage, realistically

4
major platforms closed: Instagram, Facebook, TikTok, LinkedIn

The short version. Claude Code can write almost every feature of a commercial social-listening tool. Those vendors' real advantage is their data: years of crawler infrastructure and licensed data deals. You can get your own access to news, blogs, Reddit, YouTube and reviews. In practice you can't get Meta, TikTok or LinkedIn.

Feasibility by layer



1 What's possible

A. Data collection · the hard layer

SOURCE	HOW	CAVEATS	WORTH IT?
News sites	Google News RSS, GDELT (free, global), paid news APIs	Google News RSS is unofficial and can break. GDELT is noisy. Paywalled articles give only headline and snippet. Commercial news APIs cost real money.	✓ Yes Core source, cheap, high PR value
Blogs & websites	RSS feeds, a web search API (e.g. Brave), your own crawler for a curated list	You won't match a huge crawled index. A curated list of 200–2,000 sites is realistic. Bing's search API was retired in 2025.	✓ Yes If you curate the list for your niche
Reddit	Official API	Free tier is non-commercial with rate limits. Commercial use needs Reddit's terms or a paid agreement. Deleted content must be removed from your data.	✓ Yes High-signal for B2B and tech. Check licensing first.
YouTube	YouTube Data API v3	10,000 units/day by default and a search costs 100, so about 100 searches/day. You get titles, descriptions and comments, but not what's said in the video without transcripts.	✓ Yes Titles and comments. Spoken content only for a short list.
Hacker News	Algolia HN API (free)	Developer audience only	✓ Yes If tech or dev audiences matter

SOURCE	HOW	CAVEATS	WORTH IT?
Bluesky	Open firehose and search API	Smaller audience, but fully open	✓ Yes Cheap to add, moderate value
Mastodon	Per-server public APIs	Fragmented, with no global search	⚠ Maybe Only if your niche is there
X / Twitter	Paid API	Pricing has ranged from about $200/mo to $5,000+/mo and keeps changing. Lower tiers have tight read caps and a short search lookback.	⚠ With budget Useful volume is expensive
Threads	Official API	Some public keyword search has opened up, with tight rate limits. Check the current terms.	⚠ Test it Maybe the one usable Meta-owned source
Telegram	Telegram's own API, public channels only	Needs a user account. No global search, so you monitor channels you pick.	⚠ Niche Crypto, Eastern Europe, some news
App store reviews	Apple's public review RSS per app. Google Play through scraping libraries.	Apple returns only recent reviews. Google Play has no official API for other companies' apps, so scraping is grey under its terms.	✓ Yes If you or competitors have apps
Review sites G2, Capterra, Trustpilot, TripAdvisor	Their own APIs cover your own listings only. Competitors require scraping.	Scraping breaks their terms and they actively block it	⚠ Own only Your own reviews yes. Competitors' reviews, mostly no.
Podcasts	RSS feed, then speech-to-text, then keyword search	Transcription costs grow with hours of audio	⚠ Curated 10–30 shows: yes. Broad coverage: no.
Your own social accounts IG, FB, LinkedIn page	Official APIs for accounts you own	You see @mentions and comments on your posts only, not wider conversations	✓ Yes Engagement tracking, not listening

B. AI analysis · the easy layer, with no real limits

FEATURE	HOW	CAVEATS	WORTH IT?
Sentiment #6	Claude classifies each mention, Haiku-class for cost	Handles sarcasm better than older models but not perfectly. Spot-check about 50 items per month.	✓ Yes Likely better than most vendors'
Intent #14	Claude classifies into question, complaint, purchase intent, comparison, promo, etc.	You define the categories, which is a plus for B2B	✓ Yes Very high value for sales leads
Topic analysis #13	Embeddings, clustering, then Claude names each cluster	Clusters drift between runs, so fix the topic list once it settles	✓ Yes
Semantic search #15	Embeddings in a vector store (pgvector or SQLite-vec)	None worth mentioning	✓ Yes

FEATURE	HOW	CAVEATS	WORTH IT?
Context & word cloud #9, #29	Counts of terms that appear together, or Claude-extracted themes	None	✓ Cheap Mostly useful as a visual
AI Insights briefing #11	A scheduled Claude job writes a weekly summary	Only as good as your data coverage	✓ Yes What managers actually read
AI Brand Assistant #10	Claude plus an MCP server over your mentions database	None. Very natural to build with Claude.	✓ Yes
Spike detection #12	Rolling z-score or seasonal decomposition on daily counts, then Claude explains the spike	Needs about 4–8 weeks of history before it's reliable	✓ Yes
Dedup & relevance filtering	Embedding similarity plus a Claude relevance check	Essential because keyword matching is noisy. Commercial tools have the same problem.	✓ Required Not optional

C. AI visibility · fully possible, and arguably the best place to start

FEATURE	HOW	CAVEATS	WORTH IT?
LLM brand monitoring #16	Send a set of prompts to the OpenAI, Perplexity, Gemini and Claude APIs on a schedule, then parse mentions, rank and sentiment	API answers can differ from the consumer apps (personalization, memory, different search). Answers vary run to run, so sample each prompt 3–5 times and average.	✓ Best ROI No data access problem
Prompt & competitor gaps #17	Same data, compared across competitors	Choosing the prompts is the real work. Base them on Semrush keywords and personas.	✓ Yes Feeds straight into content planning
Cited sources tracking	Parse the citations in Perplexity and Gemini answers	Not every model returns citations	✓ Yes Shows which sites to get coverage on
Google AI Overviews	Search-results APIs (SerpApi, DataForSEO)	Paid per query	✓ Yes At a modest budget
Copilot & consumer ChatGPT app	Only by scraping the web UI	Breaks their terms and is fragile	⊘ Skip Use the APIs as a proxy

Why start here: AI visibility has no data access problem, costs little to run, is the newest part of commercial listening tools (often sold as an add-on), and plugs straight into keyword and content work you already do.

D. Metrics & analytics

FEATURE	CAVEATS	WORTH IT?
Volume, trends, share of voice #5, #7	Only as complete as your sources. Share of voice is relative, so gaps affect all brands similarly.	✓ Yes

FEATURE	CAVEATS	WORTH IT?
Reach estimate	You invent the formula: followers, views and upvotes for social, site traffic for news (your Semrush connector can provide traffic). It won't be comparable with vendors' numbers.	✓ Yes Keep it simple and documented
Influence Score & influencers #8, #26	Only for authors on sources you collect. No Instagram or TikTok influencers.	⚠ B2B only Fine for Reddit, YouTube, blogs, journalists. Weak for consumer brands.
Hashtag tracking #4	Works on Bluesky, Mastodon, YouTube, and X if you pay. Not on Instagram or TikTok, where hashtags matter most.	⚠ Low value Without those platforms
Geolocation #25	Very little source data. Mostly inferred from language or the publication's country.	⚠ Approximate
Language filter #24	Easy to detect	✓ Yes

E. Reporting, alerts & integrations · all easy

FEATURE	CAVEATS	WORTH IT?
Slack / Teams alerts #3, #30, #31	Tune the thresholds or people start ignoring the alerts	✓ Yes
Scheduled email reports #20	Needs an email-sending service	✓ Yes
PDF / Excel / PPT exports #18, #19, #21	Can be branded with your SBS brand skill	✓ Yes You already have these skills
Dashboard	Streamlit or Next.js, or an Artifact dashboard for internal use	✓ Yes
Filters, Boolean, tags, saved filters #22, #23, #27, #28	Standard app work	✓ Yes
Own API & MCP server #32, #33	Lets you query everything from Claude	✓ Yes Costs very little extra
Notion export	Fits your existing workflows	✓ Yes Natural fit for you

F. Cross-cutting caveats

GDPR

Storing posts and author names counts as processing personal data. You need a legitimate-interest basis, a retention policy and a way to handle deletion requests. Get this signed off internally if it's for SBS.

No backfill

Your history starts the day you switch it on. Most APIs only look back days or weeks.

Maintenance

APIs change, scrapers break and quotas move. Expect a few hours a month of upkeep. Claude Code makes the fixes fast, but someone has to notice the breakage.

Running costs

LLM classification is cheap per mention with small models. Paid APIs (X, news, search results) are the real cost. Many sources still cost nothing.

2 What's not possible, and why

Instagram
Own account & capped hashtags only

Facebook
CrowdTangle shut down in 2024

TikTok
Research API for academics only

LinkedIn
No post-search API, sues scrapers

WHAT	WHY NOT	CAN VENDORS DO IT?
Instagram listening public posts mentioning you	Meta's API only shows your own account, @mentions of you, and a capped hashtag search (about 30 hashtags/week). It has no keyword search, and scraping breaks Meta's terms, which Meta enforces.	⚠ Partly Through licensed data partners
Facebook listening	Public page content access is restricted. CrowdTangle was shut down in 2024, and Meta's Content Library is for approved researchers only.	⚠ Partly Licensed
TikTok listening	The Research API is limited to academic and non-profit researchers. No commercial access to search public posts.	⚠ Partly Licensed
LinkedIn listening	No public post-search API. LinkedIn sues scrapers (it won against hiQ in 2022).	⚠ Very limited
Private spaces FB groups, private accounts, Discord, WhatsApp, Slack communities, closed forums	No legal access	⊘ No Vendors can't either
A vendor-scale source index tens of millions of sources	The code isn't the obstacle. It takes years of crawler infrastructure, data licensing contracts and operations. Claude Code writes software; it doesn't sign data deals.	✓ Yes That's their business
Full, affordable X coverage	Possible in principle but cost-prohibitive. Large volumes need enterprise contracts.	✓ Yes Through a commercial agreement
Historical data before you start	APIs don't offer deep history, so you can't backfill a year	✓ Yes Years of archives
Vendor-compatible numbers reach, proprietary scores, benchmarks	Their models are proprietary and calibrated on their own data. You can build equivalents, not exact copies.	n/a
Spoken content at scale all podcasts, video audio, Reels	Transcription cost at that volume, and no access to TikTok or IG media	⚠ Partial For them too
Exact consumer-app AI answers ChatGPT with memory, Copilot	The only route is scraping the web UI, which breaks their terms and gives personalized results	⚠ Approximate They approximate as well
Competitor reviews at scale G2, Capterra, Trustpilot	Their APIs cover your own listings only. Scraping breaks terms and is blocked.	⚠ Partly

Bottom line

TIER 1 · BUILD FIRST

AI visibility tracker

Use Semrush keywords to choose the prompts. Track ChatGPT, Perplexity, Gemini, Claude and AI Overviews.

Clearly worth it: fully possible, cheap, close to your content and SEO work.

TIER 2 · BUILD NEXT

Focused listening

News, blogs, Reddit, YouTube, HN and reviews, with Claude sentiment, intent and topics, a weekly briefing and Slack alerts.

Worth it for B2B. Covers most of the conversations that matter in your market.

TIER 3 · DON'T BUILD

Closed platforms

Instagram, TikTok, LinkedIn and Facebook listening, or vendor-scale crawling.

Pay a vendor if you need these. If it offers an MCP server or API, run Claude analysis on top.

Feature numbers (#1–#33) refer to our inventory of standard social-listening features. API pricing, quotas and platform access terms change often; check the current terms before committing budget. Prepared October 2026.